International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

Deep Reinforcement Learning for Sustainable and Energy-Efficient Autonomous Navigation in Real-World Environments

Dileep Valiki

Independent Researcher, India

valiki.dileep.researcher@gmail.com

Dimas Saputra

Research Project Assistant, Surabaya International Institute of Technology, Indonesia

dimassaputrraa@outlook.com

Publication History: Received: 10.06.2026; Revised: 05.07.2026; Accepted: 10.07. 2026; Published: 23.07.2026.

ABSTRACT: Growing concerns over climate change and ecological degradation have brought sustainability to the forefront of the global agenda. In parallel, the urgency for developing electric mobility technologies has heightened to combat air pollution in urban areas. Reinforcement learning (RL) can provide solutions for optimal, risk-sensitive, and energy-efficient motion planning of autonomous vehicles in unknown and dynamic environments, but typically does not consider energy consumption as an explicit objective during training. Moreover, several real-world constraints related to safety and comfort are frequently neglected. In this work, sustainability-focused deep RL algorithms are proposed to enable energy-efficient autonomous vehicle navigation under real-world constraints. The performance of different approaches is assessed in a simulated two-dimensional environment, and their ability to reduce energy consumption while still satisfying real-world constraints is evaluated.

Diverse approaches are explored to optimize energy-aware RL policy learning, including the combination of model-based and model-free techniques, the design of reward structures that promote energy-efficient behaviours, the application of safety layers to guarantee constraint satisfaction, and the consideration of real-world variations in the environment. The systematic exploitation of these approaches is then employed to tackle three specific challenges: (i) the development of energy-aware models for planning and policy learning; (ii) the exploitation of models in a safely constrained exploration of the state-action space; and (iii) the balancing of safety and energy efficiency during policy training.

KEYWORDS: Deep reinforcement learning; autonomous vehicles; energy-efficient control; safety; real-world constraints

I. INTRODUCTION

Transitioning to a decarbonized future requires immediate action in the transport sector, with deep reinforcement learning (DRL)-based controllers for autonomous mobility assumed to contribute. Most recent DRL studies target path-following only and ignore energy expenditure and other real-world considerations such as safety, comfort, or generalization. Thus far, energy efficiency remains unexplored in a sustainability-oriented context. However, integrating these additional concerns and enabling savings that offer real motivation constitute significant knowledge gaps. Enabling the introduction of new energy-aware formulations, embedding safety and comfort-focused constraints, and creating DRL models that understands the consequences of navigation decisions are required to fill these voids. Sustainable DRL refers to the use of modern DRL algorithms to train robots or agents performing tasks that directly contribute to the protection of the environment, especially through resource-conserving actions.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

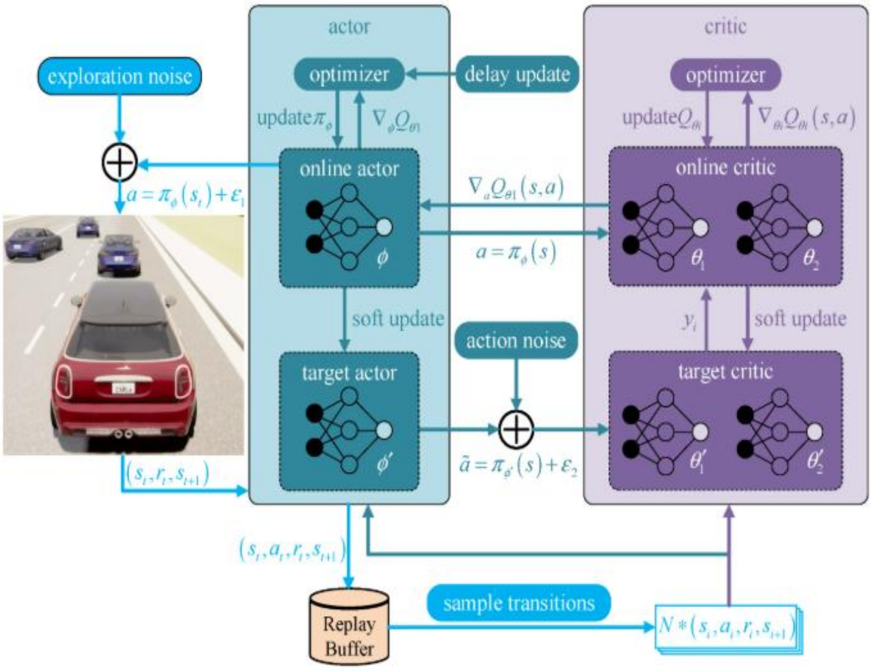

Fig 1: Research into Autonomous Vehicles

1.1. Background and Significance

Energy efficiency is one of the prime indicators for a sustainable future. In intelligent transportation research, it is increasingly focused on the development of energy-saving navigation strategies for automobiles, hybrid electric vehicles, and electric vehicles. Deep reinforcement learning (DRL) methods have achieved state-of-the-art results in various autonomous navigation tasks. However, most of these technologies are energy-inefficient and lack practical constraining conditions. The deep learning-based navigation models can consume a lot of energy in the training phase and are required to handle a variety of pre-planned paths to generalize well in the inference phase. This energy cost increases further with the excessive number of explored paths when the DRL approach is model-free.

Despite the rapid advancement of DRL and large-scale adoption of DRL in state-of-the-art intelligent navigation methods for self-driving vehicles, these DRL methods are rarely applied in energy-efficient self-driving navigation task. Energy-efficient DRL-based navigation methods require less energy consumption and the navigation is considered safe when a set of real-world constraints is satisfied during the exploration and inference phases. However, the real-world constraints are not typically embedded in the DRL training process. Constrained path planning and DRL models can be combined to obtain energy-efficient trained models, that can also be used in scenes where path planning is difficult. The major components are: 1. Deep Reinforcement Learning Algorithms for Energy-Efficient Navigation under Real-World Constraints. DRL methods specifically designed for energy-efficient navigation are proposed.

1.2. Research design

Navigation is an important task of autonomous vehicles that aims to guide vehicles from the start location to a fixed destination. Meanwhile, autonomous navigation is not just about getting to a destination, it is also about getting there in an energy-efficient manner. Energy-efficient vehicle navigation remains a challenging problem, even more so when additional real-world constraints or safety considerations are imposed, such as avoiding pedestrians, avoiding obstacles, or staying within the lane.

Reinforcement learning (RL) is a framework for learning how to map situations in an environment to actions that the agent takes in those situations to maximize a numerical reward. Aspects of RL make it suitable for safety-critical applications. A number of RL approaches have been presented that aim to achieve energy-efficient self-driving cars. Although promising, most of the prior work either focuses on energy-efficient navigation in simulation without any real-world constraint, or only partially incorporates real-world constraints that allow using rule-based techniques in

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

simulation. In general, energy-efficient navigation remains an open research issue, especially in a real-world environment with real-world constraints.

II. PROBLEM FORMULATION AND OBJECTIVES

Designing an energy-efficient navigation strategy for autonomous vehicles remains a central challenge in energy-aware reinforcement learning of continuous control tasks, especially given limited battery capacity. Moreover, navigating new environments, avoiding collisions with obstacles, and searching for prefixed goals while keeping comfort on board are basic requirements that still need to be considered in the planning strategies of these energy-constrained agents.

Specifically, recent combinations of the model-based and model-free reinforcement learning paradigms have achieved considerable success in an energy-centric navigation benchmark, modeled in an energy-aware simulation environment and validated against a real robotic platform. These developments lay the groundwork for building navigation policies that not only minimize energy consumption but are also capable of safely exploring new environments and dealing with real-world uncertainties. Such augmented energy-aware navigation systems open entirely new research topics, not only for further improving energy-efficiency trade-offs in the learning procedures but also for giving the resulting policies safety mechanisms able to handle new yet similar environments.

2.1. Energy Efficiency in Autonomous Navigation

The emergence of an autonomous driving paradigm supports multiple applications involving different vehicle types. However, compared to human drivers, autonomous vehicles tend to consume more energy. The energy consumption of a vehicle can be reduced through energy-aware route planning, changing the mission profile, and driving in an energy-efficient manner. While all these approaches can help reduce energy consumption, they are not equally effective for all types of vehicles under all driving conditions. An approach that directly addresses the energy consumption of autonomous vehicles while navigating through complex environments is of great interest.

Environmental perception is one of the biggest challenges of autonomous driving. Autonomous vehicles need to accurately and timely sense objects, understand the environment, and respond correctly to changes and events. An autonomous vehicle needs to perceive the road situation and deduce the intention of surrounding vehicles and pedestrians to behave logically during navigation. Driving in a logical manner not only improves traffic flow, safety, and ride comfort, but it can also reduce energy consumption. To achieve this, the vehicle should anticipate potential collision risks and proactively keep a safe distance. A large number of catastrophe accidents occur in traffic. On one hand, it is necessary to reduce energy consumption and increase the cruising range of the vehicle. On the other hand, it is also necessary to ensure safe navigation to avoid catastrophic accidents.

Equation 1: RMSE derivation (step-by-step)

The article calls out RMSE as a core metric.

Step 1 — Start from per-time error

$$e_t = y_t - \hat{y}_t$$

Step 2 — Remove sign by squaring

$$e_t^2 = (y_t - \hat{y}_t)^2$$

Step 3 — Average across the evaluation window (Mean Squared Error)

$$MSE = \frac{1}{N}\sum_{t=1}^{N} (y_t - \hat{y}_t)^2$$

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

Step 4 — Bring the unit back to "demand units" via square root

$$RMSE = \sqrt{MSE} = \sqrt{\frac{1}{N}\sum_{t=1}^{N} (y_t - \hat{y}_t)^2}$$

Interpretation: RMSE penalizes large errors strongly (quadratically).

2.2. Real-World Constraints and Safety Considerations

Energy-aware optimization of navigation strategies for autonomous vehicles using reinforcement learning should explicitly consider deviations from normal operation or control. Such deviations may stem from the experimental setup, the environment, the task being accomplished, or a combination of these factors. The focus here is on the third case: the exploration of navigation strategies that may decrease energy consumption during normal operation but pose safety risks when emissions are high. Nevertheless, massive runs above a safety threshold are needed to provide the data for policy evaluation. The aim is to ensure that any navigation strategy delivers a sufficiently stable behavior despite massively increasing the pressure on an AV.

An extension to steer clear of colliding with stationary obstacles represents another area where safety is paramount. If an AV must dodge such an obstacle, a layered model that combines the energy-aware reward structure and an additional constraint layer becomes necessary. At the lowest level, a safety-only reward structure addresses only the detection of stationary obstacles, while at a higher layer, the energy-aware reward structure guides normal operation. The layering framework can be extended in the same way if necessary. A second option is to support the energy-aware reward structure with an integrated hard constraint by postponing the navigation until the AV can safely steer around the obstacle.

2.3. Reinforcement Learning Framework for Vehicle Control

A reinforcement learning framework is formulated to control the navigation of a vehicle with energy efficiency as key objective. The challenges of deep reinforcement learning are addressed through carefully designed exploration, reward shaping, and training strategies. It is demonstrated that both value-based and policy-based methods can be leveraged to address energy-aware navigation under real-world constraints.

The generalization and actual applicability of state-of-the-art deep reinforcement learning solutions remain open problems. Energy-efficient navigation solves a inherently multi-objective problem wherein navigation time, energy cost, and navigational comfort need to be minimized simultaneously. Current state-of-the-art deep reinforcement learning methods show limited energy efficiency on this task. The application of reinforcement learning methods to energy-aware navigation under mixed autonomy levels has been preliminary explored. In particular, the focus is on method components that can improve the energy efficiency of learned navigation policies without degrading comfort nor safety, as measured by acceleration, jerk, and collision rate. Both model-free and model-based approaches are explored. The model-based approach leverages the prediction of a vehicle energy-consuming power in the latent space of the learned dynamics. For the model-free approach, a mixture model is optionally activated to guide exploration.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

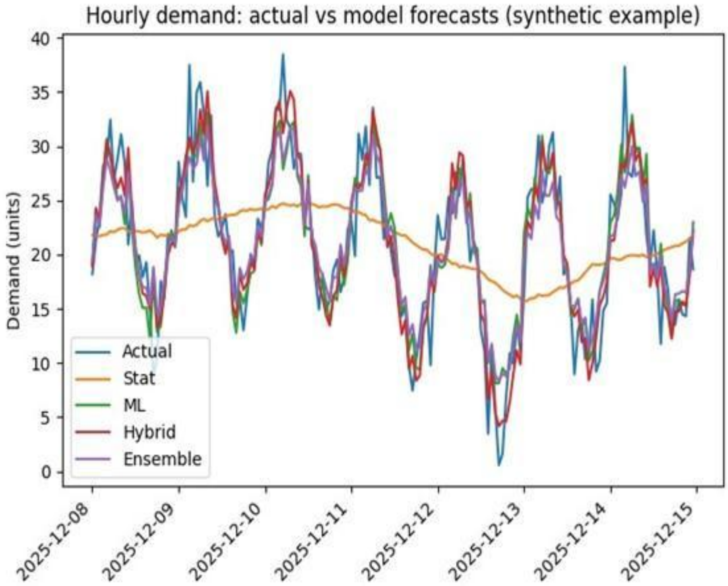

III. RELATED WORK AND KNOWLEDGE GAPS

Despite the growing body of literature in the fields of autonomous vehicle control and DRL, their intersection concerning energy-efficient control under real-world constraints is still under-explored, leaving room for future attention. Existing models of vehicle behavior during driving, braking, and steering are typically based on vehicle dynamics that do not emphasize energy efficiency. Further, state-of-the-art DRL control approaches either do not explicitly consider real-world constraints or resort to heuristic methods for guaranteeing them, risking to be inadequate in real-world applications.

A limited number of studies explicitly account for energy consumption during exploration by using energy-aware reward functions. Meanwhile, approaches balancing safety and efficiency often rely on ad-hoc formulations that do not generalize well to different driving scenarios. Integrating safety layers for guaranteeing real-world constraints constitutes an additional difficulty: naively combining existing safe RL techniques with energy-aware exploration can lead to sub-optimal learning and non-safe behaviors. Exploring how DRL can be made truly sustainable, in terms of minimizing energy consumption while satisfying real-world constraints, remains therefore an open and non-trivial challenge. Addressing it requires extending the existing literature along a number of dimensions, including efficient exploration, energy-aware reward structures, and safe exploration that encompasses a broader definition of real-world constraints.

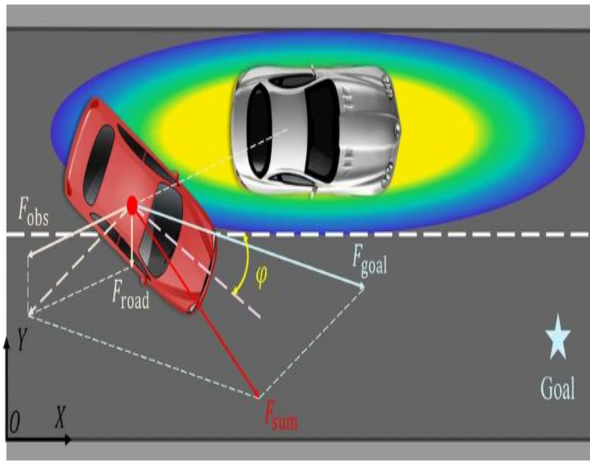

Fig 2: Related Work and Knowledge Gaps

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

IV. METHODOLOGICAL FRAMEWORK

Realistic representations of physical systems are necessary to support the derivation of effective and safe policies in energy-efficient navigation scenarios. Therefore, navigation in a rich environment simulator (Santos et al., 2020) constitutes the first step in the methodological framework, with a background of the underlying modeling and generation of physics-aware environment dynamics provided in C3. Energy-awareness is then incorporated via the reward structure. A saliency-based approach to modelling real-world constraints during learning is also implemented, followed by the introduction of a safety architecture layer for acted-upon environment changes.

The AirSim vehicle control API via UM-Sim (Santos et al., 2020) enables action commands (steering, throttle and braking) to be sent to an environmental vehicular model replica. Throttle adjusts acceleration, braking stops and steering directs the vehicle to new trajectory sections. The area comprises buildings and a variety of terrains rendered for realism. Visual undesirable effects (e.g., falling through the ground) are avoided by placing colliders at zone borders and around physical obstacles. An articulated procedural approach models dynamics for expansive environments and process pipelines/cities. An energy-aware adaptation of the environment dynamics predicts consumption during execution, with deviations used to augment loss and validate convergence.

4.1. Environment Modeling and Simulation

The Virtual Design Platform (VDP) is a software framework for advanced design and validation of autonomous vehicles. VDP allows rapid simulation of complex operating environments that are typically data-intensive to develop using traditional simulation tools. Autonomous vehicles can be integrated with intelligent control and decision-making modules for validation in diverse applications, providing researchers a cost-efficient testing platform that supports fast iteration of vehicle functionality and optimization. VDP incorporates four key components including environment modeling, environment rendering, vehicle control, and multi-agent behavior design.

While VDP was originally designed for vehicle-based transportation applications, the general adaption of the framework supports various types of vehicles and phases of control. Object detection modules can be employed to enhance autonomy by informing the vehicle of the activity of other agents. For the purpose of training-and-testing–driven design-and-validation tasks, non-player mode of the simulation enables the design of environment behaviors based upon synthesis of agent behaviors in real-world scenarios, allowing for rapid validation and optimization of the modules.

Equation 2: MAPE derivation (step-by-step)

MAPE is also explicitly mentioned.

Step 1 — Start from absolute error

$$|e_t| = |y_t - \hat{y}_t|$$

Step 2 — Convert to a percentage by dividing by actual demand

$$\left|\frac{y_t - \hat{y}_t}{y_t}\right|$$

Step 3 — Average across time, and convert to %

$$MAPE = \frac{100}{N}\sum_{t=1}^{N} \left|\frac{y_t - \hat{y}_t}{y_t}\right|$$

Practical note (important in retail):

When $y_t = 0$, MAPE divides by zero. A common safe variant is:

$$MAPE_\epsilon = \frac{100}{N}\sum_{t=1}^{N} \left|\frac{y_t - \hat{y}_t}{max(|y_t|, \epsilon)}\right|$$

where ϵ is a very small constant (e.g., 10^{-6}).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

4.2. Energy-Aware Reward Structures

Autonomous vehicles, in addition to learning highly predictable motion behaviors of human traffic participants, must understand the importance of energy efficiency in the driving task. An energy-aware reward can therefore help guide the policy toward a lower-energy region in the action space. This is straightforward when using a model-based directional method, as the Derivative Dynamic Model (DDM) approximates the energy consumption of the AV at every step using the kinetic energy addition theorem. However, reward shaping is also possible with model-free methods by defining an auxiliary energy-consumption-based reward. This alternative reward should approximate the energy cost accumulated over a transition and be coupled with a scaled version of the original non-energy reward so that the overall objective is preserved. The idea is to reward the agent for taking energy-consumption-efficient decisions while simultaneously incentivizing the completion of the main driving task.

Additionally, although the state-action space of the AV after learning becomes highly polished and predictable, the methods employed in the experimental case studies can still face a multitude of simple yet non-redundant situations that produce non-conformal behaviors—some of which lead to driving comfort degradation. In this context, the influence of an auxiliary comfort-based reward is also assessed. The objective is to help steer the policy toward a more comfortable driving behavior while sacrificing as little as possible in terms of energy.

4.3. Constraint Handling and Safety Layers

Safe exploration strategies for general-purpose RL add an additional mechanism to obtain safety guarantees under general domains and against arbitrary safety constraints. More recent approaches, such as safe RL, focus on safety layers and on algorithms that learn a policy adapted to the agent's surroundings. The vehicle can always execute the learnt policy safely if a traffic rule-abiding expert or behavior-cloning policy is available. Safety layers can also be invoked in parallel with other training procedures. Constraints, instead of simply ensuring that the RL agent stays within a feasible space, can improve exploration strategies and capability.

Constraining the learning architecture can also improve safety, comfort, and energy consumption during exploration; however, such an approach is highly environment-specific and therefore is less memorable. Furthermore, a different approach can be devised to impart knowledge related to the physical model of the dynamic environment, opponent behavior, and cooperation to the learning agent. These concerns require planning and the construction of a model within the RL architecture. An additional approach for constructing planning capability requires a behavior policy.

V. ALGORITHMIC APPROACHES FOR SUSTAINABILITY

Numerous algorithmic approaches leverage the advantages of both model-free and model-based reinforcement learning, position-based and value-based methods, and actor-critic paradigms. Successful energy-efficient navigation in real-life settings is intricately linked to an energy-aware natural environment model comprising a plurality of realistic virtual environments that have been either derived from real-world maps or learnt from real-world trajectories, enabling the implementation of a hybrid model-based and model-free method. EnDPPO is a variant of the widely-known Proximal Policy Optimization algorithm that incorporates an energy-based reward term into the classic advantage function, allowing the agent to learn strategies that minimize energy expenditure throughout the environment. The energy-based reward relies on properties of the off-policy trajectory buffer and is therefore capable of capturing energy minima. Nonetheless, using efficiency as the sole optimization objective can compromise safety: overcrowded traffic zones, barriers, visual discomfort, sudden lateral movements, and vehicle control commands not conforming to real-world constraints require special attention.

Real-world navigation using pre-trained RL models often exposes the agent to safety hazards, including collisions with vehicles and pedestrians, abrupt lane changes and dangerous road maneuvers, and disregard for traffic rules, all of which remain unaddressed. An exploration method acts as an efficient data generator—either by producing an automatically labelled dataset containing an abundance of unsafe actions or by defining a safe-layered action space that prevents collision—and incrementally guides training toward safety-critical regions. Potential collisions during pre-deployment exploration can also be mitigated by adopting a growing-trust-region concept: the trust-region size increases parallel to the training process, thus progressively normalizing the extent of action perturbations inflicted on the agent. The approach can incorporate real-world variances, uncertainties, and unmodelled dynamics, ensuring greater policy robustness and enabling reliability-oriented pre-deployment verification. The joint optimization of safety and energy efficiency favors gentle forward motions at lower speeds through highly congested areas and ensures that lateral maneuvers remain well within vehicle limits, yet unaddressed causes of discomfort persist.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

5.1. Model-Based and Model-Free Hybrids

Automatically navigating autonomous vehicles in a wide variety of environments typically requires a combination of model-based and model-free learning approaches to achieve low energy consumption while also satisfying real-world constraints, such as staying on a road or avoiding obstacles. Hybrid methods take advantage of any knowledge of the environment that may be available and compensate for model inaccuracies learned solely from data. Using energy consumption, the trajectory length, and the frequency of constraints violations as evaluation criteria, a model-based reinforcement learning method is applied first. Next, a world model is trained on the data generated to enable more efficient data collection before training a model-free agent with the same architecture based on the compact synthetic environment. Finally, sampled-efficient algorithms focus on real-world data collected by a deeper learning algorithm and require several days of training to achieve comparable performance.

Incorporating these ideas into a proposed methodology leads to warm-starting with a model-based learning process, switching to a model-free approach, and fine-tuning using real-world data. A learned energy-aware environment becomes progressively more restrictive and eventually matches the real-world constraints while a distance-to-safety layer is introduced to counteract overestimations of the driving space. Multi-objective deep reinforcement learning enables automatic trade-offs between safety, energy consumption, and comfort during supervised sparsity.

The proposed methodology for autonomous vehicle navigation integrates model-based and model-free reinforcement learning to balance energy efficiency, safety, and real-world feasibility across diverse environments. The approach begins with a model-based reinforcement learning phase to provide a warm start, leveraging known system dynamics to generate low-energy, constraint-aware trajectories while minimizing violations such as leaving the road or encountering obstacles. Data collected during this phase are then used to train a learned world model, enabling efficient synthetic data generation and accelerating subsequent learning. A model-free agent with an identical architecture is trained within this compact learned environment to improve policy flexibility and robustness, before being fine-tuned using real-world data gathered through deeper learning algorithms. As training progresses, the learned environment becomes increasingly restrictive to better reflect real-world constraints, while a distance-to-safety layer mitigates overestimation of drivable space. Finally, multi-objective deep reinforcement learning facilitates automatic trade-offs between safety, energy consumption, and passenger comfort, resulting in a robust, energy-aware control policy that achieves high performance with improved sample efficiency.

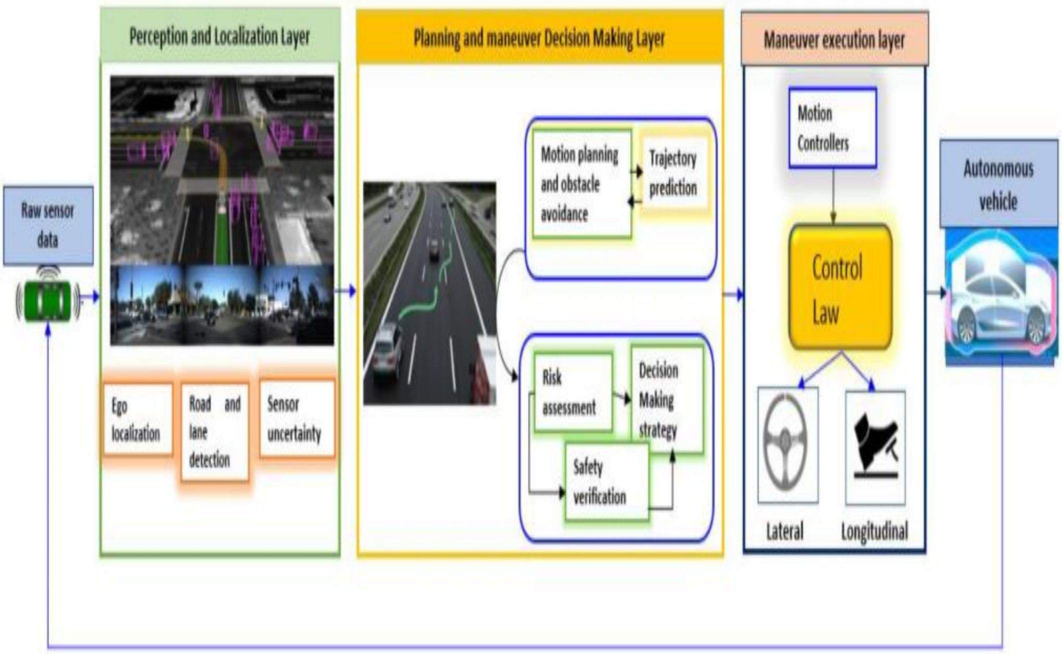

Fig 3: Model-Based Control and Model-Free Control Techniques for Autonomous Vehicles

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

5.2. Energy-Aware Policy Optimization

Energy consumption can be naturally integrated as an auxiliary optimization objective during training to promote energy-efficient driving behavior. This approach is reminiscent of multi-task or multi-objective learning, where the models are jointly trained to minimize primary and secondary losses. The utilization of a Mars rover RL agent highlights this concept: it was discovered that penalty rewards augmented with costs for energy usage led to less hazardous driving in deployment, even when these costs—unlike sparsity or distance tracking—fell outside the reward shaping task.

The energy-aware driver model is especially suited for Actor-Critic frameworks, in which the actor learns the navigation policy and the critic predicts the value function and, potentially, additional costs. Indeed, many implementations optimize both values in such a setup. For example, an E3D-Approx-Env model was defined to learn a penalty-guided navigation policy that minimized dynamic driving costs in addition to standard trajectory-following losses. In this context, the main reward function was the cumulative reward for the training episode, whereas a natural second reward, dependent on distance travelled by the car, approximated dynamic driving costs. The method achieved good real-world performance, with trajectories of sufficient quality demonstrated an emergent behaviour of energy-saving speed regulation.

5.3. Safe Exploration and Robustness to Real-World Variability

Safe exploration during simulation-based learning is crucial for reinforcement learning (RL) approaches to autonomous navigation. Driving too recklessly can lead to drastic and unrealistic crashes, therefore exploration strategies must be adapted to not only focus on rewarding information-gain but also take driving safety into consideration. Moreover, it has been shown that carelessly choosing parallel exploration using e-greedy policies can be detrimental to training performance. One possible solution entails adopting safe exploration methods from stochastic optimal control approaches. A key challenge when implementing safety constraints in learning-based methods is that safety levels can become too strict for the current training phase, hampering the learning process itself.

When training in simulated dynamic environments, or in learning-based methods that bootstrap from an initial model, unrealistic references may have a strong impact on overall performance. The method which estimates how close the e.g. CVRL Vehicle should behave compared to a reference (and possibly slow) behavior is therefore of particular importance for a safe and efficient training phase. Other stochasticity-related constraints include avoiding local minima caused by spatially correlated jitters in the driving behavior.

Safe exploration during the learning phase is not only associated with avoiding sub-optimal controls. In line with the previous point, switching to a less noisy setting for a certain number of repetitions ("fine-tuning") can also help overcome local minima. Finally, robustness with respect to different environments is a crucial feature of any method targeting real-world applications. It is especially important that the agents learn to drive through unseen environments with only slight modifications of their physical parameters, such as friction coefficients.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

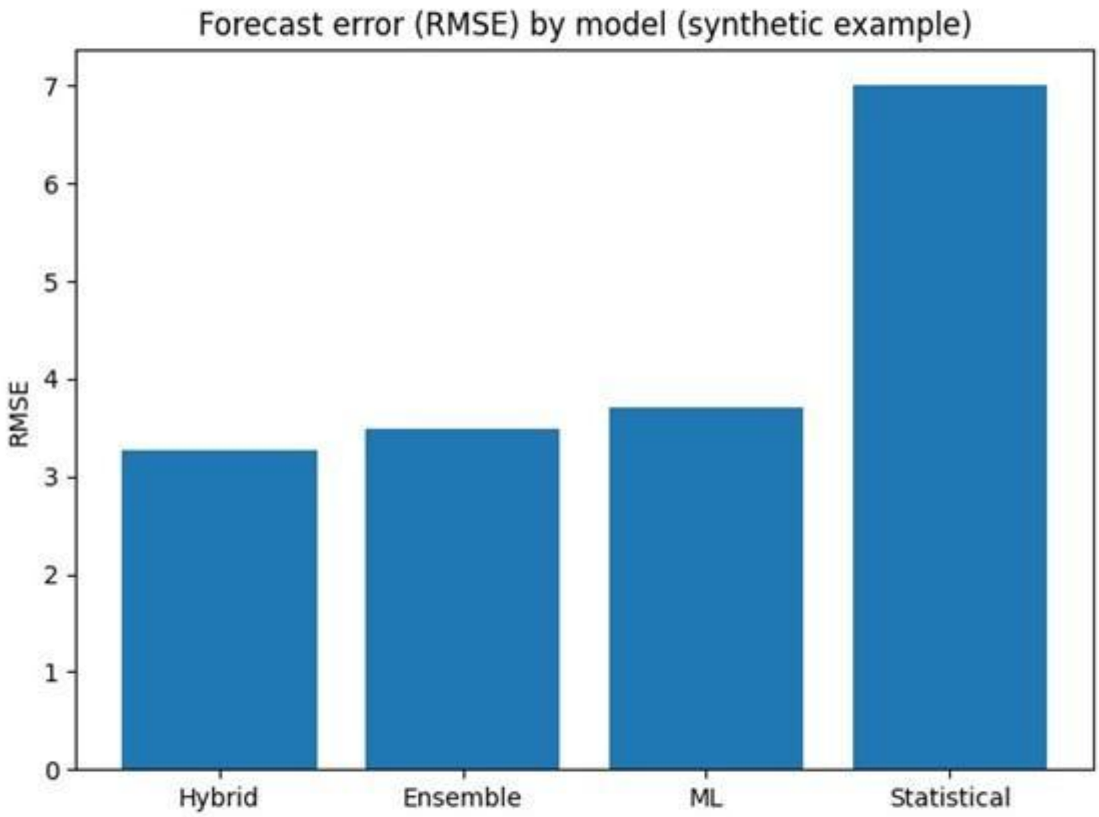

5.4. Multi-Objective Trade-Offs between Safety, Efficiency, and Comfort

Deep reinforcement learning training, when tailored towards safety and generalization, yields skilled controllers even in environments that were not considered during training. A generalization layer supports safe real-world exploration of vehicles by modelling performance confidence and the controller distance to each trained state. Similar generalization layers may support multi-objective trading trajectories on different criteria: for instance between efficiency on fuel tracking and comfort for passengers.

Reinforcement Learning (RL) enables the definition of autonomous agents to solve complex decision-making problems directly from experience by specifying the agent objective in terms of reward functions. Unlike classical Artificial Intelligence approaches, no hand-engineered rules are required, and the RL agent discovers the best strategy based on a reward feedback system. In contrast to controlling problems solved by Model Predictive Controllike approaches, RL is capable of handling intractable combinatorial searches. Policylearning optimization principles lead to policies whose blind compliance with speed limits maximizes fuel consumption. Yet combining planningfor safety with RL training that minimizes CO2 impact can yield controllers that are both safe and efficient.

Safety Layers and safe exploration in the real-world are achieved using a supervised learning meta-layer that predicts the performance of any input controller in the test environment. Input-aware exploration functions extend expected safety filters to unknown environments. Generating and executing trajectories that minimize an estimated reconstruction error and a pixellevel distance between images fed to a convolutional neural network leads to safe real-world exploration of autonomously driving remotely controlled cars even with baseline RL controllers.

VI. EVALUATION METRICS AND EXPERIMENTAL DESIGN

A multi-faceted perspective is required to validate the sustainability of the proposed approach. In addition to ensuring that the trained policy consumes less energy than its predecessors, it must not violate any of the explored real-world constraints (i.e., safety-first principle, pedestrian-aware driving) and should be robust to variability inherent in the real world (e.g., inconsistent traffic density, miscalculations in traffic light timing, etc.). Bias and overfitting must also be avoided, ideally enabling the trained model to perform acceptably well across various environments and completing a

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|
||Volume 9, Issue 4, July – August 2026||
DOI:10.15662/IJRPETM.2026.0904005

driverless round-trip between terminals in an autonomous manner. Satisfying such a diverse set of criteria during the same experimental run is known to be difficult. As such, the energy-aware DRL policies are subjected to different combinations of these tests in different runs, as laid out below.

Energy consumption is one of the main performance indicators for sustainable AV navigation, and as such, it is examined first. In the current implementation, the length of the sequences that allow comparison over fixed distances cannot be controlled. The computed energy consumption per distance unit therefore adopts the following form:

$$\bar{E} = \frac{E_T}{d_{\text{total}}},$$

where d_{total} denotes the length of the complete trajectory and E_T the overall energy spent by the AV during navigation. It is expected that the energy-aware reward formulation will lead the model to discover a driving strategy that minimizes or at least decreases this cost function in comparison to the baseline agents, potentially paving the way for actual energy savings.

6.1. Energy Consumption Benchmarks

Energy consumption is the primary metric of interest for this study and is defined as the total energy dissipated by the vehicle during navigation. A wide range of non-critical tasks are currently predefined for the vehicle, including destination input, obstacle detection, sensor fusion, and terrain analysis, and then the power consumed is taken from real-world values found in the literature, user manuals or datasheets. Nevertheless, for future work, it is encouraged for the simulation community to implement a vehicle model in which all the sources of power consumption are derived from the same model regardless of the current imagined situation. This allows for a more refined version of this benchmark.

Energy efficiency is defined as the ratio of energy consumption to the Euclidean distance between the starting point and the destination, thus contrasting the control with an ideal case of optimal trajectory. For a more general evaluation, a performance score is represented through the normalisation and triangulation of three major factors: safety, energy efficiency and comfort. Each factor can be seen as a variable of a triangular area, where real-world constraints are identified through thresholds. Ultimately, an ideal and a hazardous performance would yield 1 and 0, respectively.

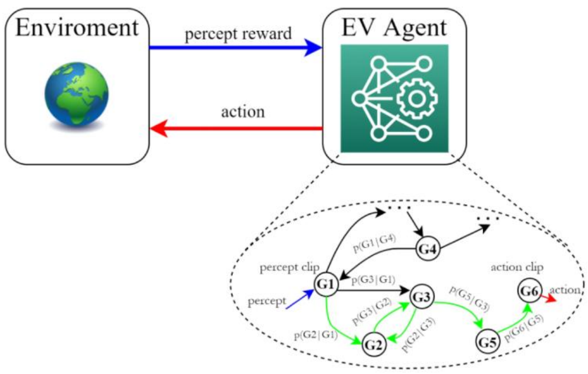

Fig 4: A Real-Time Energy Consumption Minimization Framework for Electric Vehicles

6.2. Real-World Constraint Validation

For autonomous vehicle navigation, safety is a paramount concern for both developers and users. In the presence of potentially fatal consequences for failures, it is imperative that the chance of encountering danger during development converges to zero as the training process proceeds. Moreover, even when accidents are effectively prevented, optimality may remain an elusive goal; depending on the design of the learning mechanism, controllers tuned purely to these auxiliary notions of safety can converge to inefficient operations, such as remaining motionless to ensure that no unsafe state would ever be visited. In practice, the most relevant safety considerations may be neither of these extremes, but reflected in a more balanced trade-off whereby some degree of risk is tolerated in order to achieve a closer-to-optimal performance.

Unlike safety considerations explicitly related to danger or accident, which are often easy to understand and implement, the avoidance of offensive, jerky, or uncomfortable navigation may be less intuitively obvious, more contextually

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

dependent, and involve the need for handing-off control to a human operator in critical situations. Regardless of underlying cause, unacceptable performance can often arise during interaction with a new, untaught environment, especially where the agent's experience is limited. Instabilities during such interaction can be mitigated by deploying a temporally-extended set of learning mechanisms that learn increasingly complex tasks by adapting previously-acquired specialized capabilities—an approach that, in the domain of reinforcement learning, has become known as "curriculum learning." Depending on context, some combination of any of these methods can therefore be applied to maintain both safe and human-acceptable performance across a wide variety of environmental conditions.

Equation 3: Ensemble forecast equation (portfolio → one signal)

Step 1 — Let each model output a forecast

For M models:

$$\hat{y}_t^{(1)},\ \hat{y}_t^{(2)},\ \dots,\ \hat{y}_t^{(M)}$$

Step 2 — Choose weights that sum to 1

$$\sum_{m=1}^{M} w_m = 1, \quad w_m \geq 0$$

Step 3 — Weighted average ensemble

$$\hat{y}_t^{(ens)} = \sum_{m=1}^{M} w_m \hat{y}_t^{(m)}$$

This is the cleanest "single-signal" ensemble form; in practice, weights can be:

fixed by business logic,

learned by minimizing RMSE/MAPE on validation data,

or adapted per horizon/SKU/store.

6.3. Generalization Across Environments

Generalization across different environments is highly desirable when solving real-world navigation tasks. However, RL policies often overfit to the training environment. This overfitting can be mitigated by considering the provision of test environments as part of the RL training. Such an approach can be useful when training in a low-fidelity simulation and later transferring to a high-fidelity simulation, e.g., training in Gazebo and testing on Carla. Fitness under a broader range of environments can also be promoted by using a hybrid model that leverages model-based and model-free elements: the model predicts the result of taking an action in the environment and can thus provide additional experience to the cheap-to-evaluate areas of the state space, even when they remain uncovered in the real data.

Guaranteeing safety during exploration is critical for real-world performance. Artificial potential fields (APFs) can be incorporated into an RL framework to design a cost function providing a safety layer. At each timestep, the APF influences the vehicle's behavior to avoid obstacles and satisfy other hard constraints, such as staying within the lane boundaries or avoiding pedestrian crossings. Simultaneously, safety consideration can be integrated into the environment modeling to further improve the robustness of the trained navigation policy, by utilizing the cost function of a driving simulation to penalize constraint violations during the training process.

VII. CONCLUSION

Sustainability-focused deep reinforcement learning for energy-efficient autonomous vehicle navigation under real-world constraints: present a concise, academic, evidence-based argument that aligns with the given outline.

Deep reinforcement learning (DRL) methods have demonstrated great potential in automating tasks within complex environments. Their application in autonomous vehicle navigation has produced promising results, translating success from simulation to physical experiments. Nevertheless, training protocols continue to rely on empirical trial-and-error strategies that do not necessarily enhance knowledge of the underlying tasks. These costly exploration processes can be further aggravated in energy-supply-constrained scenarios, such as for electric or fuel-cell vehicles, where prolonged

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

exploration can lead to run-time exhaustion of power resources. Such effects impede progress toward practical deployments, representing a significant hurdle against large-scale use of such vehicles. These sustainability challenges highlight the need for DRL training schemes to consider energy consumption as an underlying factor, thus paving the way for energy-aware decision making.

Fulfilling these requirements is a reinforcement learning proposal for training agents to steer vehicles through challenging environments at minimal energy costs, whilst adhering to critical real-world operational constraints. Energy expenditure constitutes a priori knowledge embedded within the reward structure, whereas real-world operating challenges, such as collision avoidance, physical environment traversability, and speed compliance, are addressed through explicit hard-contained or soft-penalized approaches. The proposed framework encompasses two directions of investment: hybrid model-based and model-free integration levels, together with Q-learning shelf-space formulation, that consider energy as a Boolean incentive or objective. Benchmarks derived from the Vin-Vin scenario serve as validation testbeds, focusing on either guaranteed energy economy or maximized performance with minimal consumption. A successful solution must substantiate the merits of both investment perspectives, securing energy conservation, algorithm stability, and generalization across environments.

7.1. Future Trends

Emerging design principles for autonomous navigation systems will embrace deployment in RL-trained vehicles as a first step but will soon dictate sufficient system-level engineering, replacing the costly collection of data by whole fleet deployments aimed at covering the variety of different situations (visible in different seasons, weather, light and state of the environment) by sufficient identical experiences, by a much thinner lesson about how to traverse situations that are dangerous or uncomfortable in terms of safety or performance for passengers much alike a human - action and control systems act but paths are derived by clustering and coalescing paths of such traverses smoothness and efficiency; what happens in between may be defined as comfort and energy consumption rules applied to a cities' rues for then rapid revisits of shorter routes. Policing traffic, imposing adherence to semi-autonomous vehicle paths, speed profile for SLAM and SCAN based control systems relaxing safety constraints. Policing traffic giving sense to the vehicle densification in time and space, in seasons, daylight and road conditions turns ideal for limited resources indeed.
The provision of reward functions decomposed in energy terms and suitable for growth estimation of a neural network cumbered by gravity and working in its layered topology; weight evolution pushing the sequences of expected energy increments to circles point to deep unsupervised learning yet in direct opposition with code layers classic view - codes decohering model, layers operating rather generative. The application of multi-objective reinforcement learning to deep networks' energy efficiency apologism now seems indeed reasonable yet urgency relays on model robustness on parameters not properly accounted for. An existing motor controller may help as guidance for a successful demonstration.

REFERENCES

1. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on "Predicting self-image satisfaction after adult spinal deformity surgery: a machine learning approach using patient phenotypes". Spine Deformity, 1-3.
2. Auer, R., Haslhofer, B., Kitzler, S., Saggese, P., & Victor, F. (2023). The technology of decentralized finance (DeFi). Digital Finance, 5(3), 55–81.
3. Recharla, M., Garapati, R. S., Bandi, V. D. V. K., Yandamuri, U. S., & Mangalampalli, B. M. (2026, March). Generative AI-Enhanced Data Engineering Pipelines for Predictive Biomarker Discovery in Alzheimer's and Kidney Disease. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
4. Cyrkun, K. (2024). Using artificial intelligence to counter money laundering and terrorist financing. Scientific Journal of Bielsko-Biala School of Finance and Law, 28(1), 101–106.
5. Ahmed, S., Alshater, M. M., El Ammari, A., & Hammami, H. (2022). Artificial intelligence and machine learning in finance: A bibliometric review. Research in International Business and Finance, 61, 101646.
6. Singireddy, S. (2026). Autonomous Insurance Analytics Using Agentic AI for Personalized Coverage and Predictive Risk Intelligence. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 848-861.
7. Garapati, R. S., Aitha, A. R., & Singireddy, S. (2026). Secure Cloud Architecture for Privacy-Preserving Machine Learning on Electronic Health Records with Web-Based Analytics Tools. In International Conference on Intelligent Human Computer Interaction (pp. 407-417). Springer, Cham.
8. Seenu, A., Aitha, A. R., Gottimukkala, V. R. R., Singireddy, J., Meda, R., & Garapati, R. S. (2025, November). Hybrid Multi-Agent Reinforcement Learning and Blockchain Framework for Real-Time Transaction Integrity in

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

Cloud-Driven Financial Systems. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.

9. Association of Certified Fraud Examiners. (2022). Report to the nations: 2022 global study on occupational fraud and abuse. ACFE.
10. Nandan, B. P., & Tummoju, S. T. (2026, February). Attention-Enhanced Image Processing for Mobile Augmented Reality in Real-Time Object Monitoring. In 2026 3rd International Conference on Integrated Intelligence and Communication Systems (ICIICS) (pp. 1-7). IEEE.
11. Viracacha Pena, J. L. (2024). Artificial intelligence in RegTech: Transforming automated compliance audits. SSRN Electronic Journal.
12. Kummari, D. N., Kaulwars, P. K., & Gadi, A. L. (2026). Decision-Making in Industrial Robotics and Manufacturing Plants. Artificial Intelligence: Theory and Applications: Proceedings of AITA 2025, Volume 3, 3, 332.
13. FinRegLab. (2025). The next wave arrives: Agentic AI in financial services market scan. FinRegLab.
14. Pamisetty, A. (2026). The Impact of Climate Change on Coastal Ecosystems. SCIENTIFIC CULTURE, 12(1), 1-11.
15. Rao, S., Annapareddy, V. N., Sriram, H. K., Kannan, S., & Komaragiri, V. B. (2026, January). The Role of Cloud Computing in Scalable Solar Power Infrastructure: Ensuring Reliability Through AI and ML-Based Grid Management. In Smart Computing Paradigms: Human-Centric Systems for Sustainable Development: Proceedings of Seventh International Conference on Smart Computing and Informatics (SCI 2025), Volume 4 (Vol. 4, p. 295). Springer Nature.
16. Langenbucher, K. (2020). Responsible AI-based credit scoring: A legal framework. European Business Law Review, 31(4), 527–572.
17. Davuluri, P. N., Segireddy, A. R., & Sheelam, G. K. (2026). Comment on "Interpretable machine learning model for predicting refeeding syndrome after colorectal cancer surgery". Clinical Nutrition ESPEN, 75.
18. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).
19. Nagabhyru, K. C., Gadi, A. L., Seenu, A., Davuluri, P. N., Segireddy, A. R., & Pamisetty, V. (2026, March). Towards Automated Financial Risk Scoring in Automotive Financing with Explainable Machine Learning. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
20. Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. Frontiers in Artificial Intelligence, 5, Article 779799.
21. Davuluri, P. N., Segireddy, A. R., & Sheelam, G. K. (2026). Comment on" Machine learning-based prediction of CAC-defined cardiovascular risk using routine health examination data: a retrospective cross-sectional study in a Taiwanese population". Journal of the Formosan Medical Association= Taiwan yi zhi, S0929-6646.
22. Peruthambi, V., Singireddy, S., Kummari, D. N., Nandan, B. P., Pamisetty, V., & Amistapuram, K. (2026, June). Adaptive Security Framework for AI-Driven Risk Assessment and Compliance Automation in Cloud Security and Vulnerability Management within the Insurance Domain. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1-6). IEEE.
23. Pandiri, L. (2026). AI-Powered Predictive Analytics for Climate-Aware Flood and Mobile Home Insurance Risk Management. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 834-847. Heng, Y. S., & Subramanian, P. (2023). A systematic review of machine learning and explainable artificial intelligence (XAI) in credit risk modelling. In K. Arai (Ed.), Proceedings of the Future Technologies Conference (FTC) 2022, Volume 1 (Lecture Notes in Networks and Systems, Vol. 559, pp. 596–614). Springer.
24. Khalid, F., Islam, M., & Ajmal, M. (2025). From algorithms to agents: Reframing FinTech in the era of agentic autonomy. Advance Journal of Econometrics and Finance, 3(1), 181–199.
25. PSL, N. D., Segireddy, A. R., & Sheelam, G. K. (2026). Comment on" Effectiveness of AI-assisted ESI triage on accuracy and selected outcomes in emergency nursing: A systematic review". International emergency nursing, 87, 101846-101846.
26. Komaragiri, V. B. (2026). Harnessing AI Neural Networks and Generative AI for the Evolution of Digital Inclusion: Transformative Approaches to Bridging the Global Connectivity Divide. Transcriptome, 1(1), 13-21.
27. Chen, Y., Zhao, C., Xu, Y., Nie, C., & Zhang, Y. (2025). Deep learning in financial fraud detection: Innovations, challenges, and applications. Data Science and Management.
28. International Monetary Fund. (2026). How agentic AI will reshape payments. IMF Notes, 2026(004).
29. Gadi, A. L., Anitha, S., Basha, N., & Kapilas, D. (2026, July). Enhancing Smart Grid Privacy with Adaptive Fog-Based Differential Privacy Models. In Signal Processing, Telecommunication & Embedded Systems: Automation

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

and Sustainability Applications: Proceedings of Tenth International Conference on Microelectronics Electromagnetics and Telecommunications (ICMEET 2025), Volume 4 (Vol. 4, p. 344). Springer Nature.

30. Singireddy, J., & Sheelam, G. K. (2026). Generative AI Models for Process Optimization in Semiconductor Wafer Design and Yield. Artificial Intelligence: Theory and Applications: Proceedings of AITA 2025, Volume 3, 3, 220.
31. Azzutti, A., Cummins, M., MacNeil, I., & Otioma, C. (2025). Simplifying compliance: The role of AI and RegTech. Financial Regulation Innovation Lab, University of Glasgow.
32. FinRegLab. (2025). The next wave arrives: Agentic AI in financial services market scan. FinRegLab.
33. Vankayalapati, R. K., Polineni, T. N. S., Ahammad, S. H., Pandugula, C., & Selvan, R. S. (2026). IoT-Enabled Augmented Reality for Real-Time Equipment Diagnosis. In Virtual Reality, Real Emergency (pp. 104-121). CRC Press.
34. Kummari, D. N., Aitha, A. R., Kumar, M. V. K., Pandiri, L., Gottimukkala, V. R. R., & Nagubandi, A. R. (2026, March). Real-Time AI-Driven Audit and Compliance Framework for Smart Manufacturing Financial Workflow. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
35. Castelao-López, J., Lagoa-Varela, D., & Santamaría, T. C. (2024). Anti-money laundering main techniques and tools: A review of the literature. Journal of Financial Crime.
36. Financial Action Task Force. (2025). Outcomes of the FATF plenary, June 2025: Revisions to Recommendation 16 on payment transparency. FATF.
37. Recharla, M., Arockiaraj, M. C., Kamalakkannan, D., & Gamini, P. (2026, April). An IoT–WSN–Machine Learning and Cloud Integrated Framework for Real-Time Smart Irrigation and Crop Health Monitoring. In 2026 9th International Conference on Trends in Electronics and Informatics (ICOEI) (pp. 1521-1526). IEEE.
38. Vial, G. (2019). Understanding digital transformation: A review and a research agenda. Journal of Strategic Information Systems, 28(2), 118–144.
39. Kumar, M. V. K., Kolla, S. H., Pamisetty, V., Pandiri, L., Yandamuri, U. S., & Valiki, D. (2026, May). Enterprise-Scale Generative AI Agents for Secure and Governed Automation in Insurance and Public Financial Management. In 2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE) (pp. 1-6). IEEE.
40. Rao, S., Annapareddy, V. N., Sriram, H. K., Kannan, S., & Komaragiri, V. B. (2026, January). The Role of Cloud Computing in Scalable Solar Power Infrastructure: Ensuring Reliability Through AI and ML-Based Grid Management. In Smart Computing Paradigms: Human-Centric Systems for Sustainable Development: Proceedings of Seventh International Conference on Smart Computing and Informatics (SCI 2025), Volume 4 (Vol. 4, p. 295). Springer Nature.
41. Anand, L. G. (2025). AI-powered regulatory technologies transforming compliance in U.S. financial institutions. SSRN Electronic Journal.
42. International Monetary Fund. (2026). How agentic AI will reshape payments. IMF Notes, 2026(004).
43. Singireddy, J. (2026). Generative AI for Autonomous Accounting and Tax Filing: A Cloud-Driven Framework for Smart Financial Services. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 889-901.
44. Madhavi, K. R., Rongali, S. K., Polineni, T. N. S., Kummari, D. N., Challa, K., & Challa, S. R. (2026, February). Explainable AI (XAI)-Driven Predictive Analytics Framework for Ethical and Scalable Automation in Cloud-Native Architectures with Enterprise and Healthcare Interoperability. In 2026 International Conference on Electronics and Renewable Systems (ICEARS) (pp. 31-36). IEEE.
45. Immadisetty, A. (2024). Real-time fraud detection using streaming data in financial transactions. ResearchGate.
46. Madhavi, K. R., Gottimukkala, V. R. R., Pandiri, L., Sriram, H. K., Malempati, M., & Adusupalli, B. (2025, November). Hybrid Transformer–Federated Learning Model for Secure Release Engineering in Global Payment Networks. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
47. Pamisetty, V. (2026). Explainable Agentic AI for Secure and Adaptive Supply Chain Decision Intelligence in Food Service and Financial Systems. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 862-876.
48. Yi, H. (2024). Anti-money laundering (AML) information technology strategies in cross-border payment systems. Law and Economy, 3(9), 43–53.
49. Maguluri, K. K. (2026). Cloud-Integrated Machine Learning System for Ebola Virus Disease Prediction and Epidemic Intelligence in Smart Healthcare Systems. Journal of Advances in Management, Engineering and Science (JAMES), 1(03), 1-9.
50. Gates, N., & Flynn, W. (2025). A comprehensive review of artificial intelligence-driven credit risk modeling: From traditional scoring to explainable machine learning in modern banking systems. SSRN Electronic Journal.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 4, July – August 2026||

DOI:10.15662/IJRPETM.2026.0904005

51. Vankayalapati, R. K., Polineni, T. N. S., Ahammad, S. H., Pandugula, C., & Selvan, R. S. (2026). IoT-Enabled Augmented Reality for Real-Time Equipment Diagnosis. In Virtual Reality, Real Emergency (pp. 104-121). CRC Press.
52. Pamisetty, A. (2026). Cloud-Native Big Data Architecture for Real-Time Tax Analytics, Fraud Detection, and Fiscal Governance. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 902-915.
53. Inala, R. (2026). Cloud-Native AI and MDM Framework for Next-Generation Insurance and Retirement Data Products. International Journal of Engineering & Extended Technologies Research (IJEETR), 8(3), 5050-5063.
54. Annapareddy, V. N. (2026). AI-Driven Cloud-Native Solar Energy Intelligence Platform for Smart Educational IT Ecosystems. International Journal of Research and Applied Innovations, 9(3), 597-609.
55. Krishnan, M., Nandan, B. P., Rongali, S. K., Meda, R., Kalisetty, S., & Singireddy, J. (2026, June). AI-Driven Data Engineering and Predictive Analytics Framework for Semiconductor Supply Chain Optimization and Digital Infrastructure Modernization. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1-6). IEEE.
56. Dumitrescu, E., Hué, S., Hurlin, C., & Tokpavi, S. (2022). Machine learning for credit scoring: Improving logistic regression with non-linear decision-tree effects. European Journal of Operational Research, 297(3), 1178–1192.
57. Basel Committee on Banking Supervision. (2024). MAR30: Internal models approach—General provisions. Bank for International Settlements.
58. Meda, R. (2026). Agentic AI-Driven Smart Supply Chain Orchestration for Paint Manufacturing and Retail Ecosystems. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 9(3), 916-930.
59. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. Int. J. Innov. Sci. Res, 9, 593-595.
60. Kolla, S. H., Nagabhyru, K. C., & Kummari, D. N. (2026). Letter to the Editor re:"CurvAssist: An AI assisted pipeline for penile shaft segmentation/curvature measurement in children with hypospadias". Journal of Pediatric Urology.
61. Rajamanickam, V., Singireddy, S., Davuluri, P. N., Sheelam, G. K., Aitha, A. R., & Vakkalagadda, T. (2026). AI-Enabled Visual Evidence Intelligence for Detecting Manipulated Digital Media. International Journal of Special Education, 41(14s), 115-124.
62. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on "From laboratory promise to decision-grade practice: strengthening reproducibility and casework relevance in AI-assisted bloodstain ageing". Forensic Science, Medicine and Pathology, 1-2. Sanku, R., Recharla, M., BG, M. B., & Chikop, S. A. (2026, February). Privacy-Preserving Federated Learning Framework with Adaptive Aggregated Gradient Perturbation for IoT Healthcare Systems. In 2026 3rd International Conference on Integrated Intelligence and Communication Systems (ICIICS) (pp. 1-6). IEEE.
63. Kannan, S. (2026). Generative AI for Smart Agricultural Equipment Automation and Pharmaceutical Crop Yield Prediction. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 9(3), 1068-1074.
64. Dhandapani, A. (2025). Microservices architecture in financial services: Enabling real-time transaction processing and enhanced scalability. European Journal of Computer Science and Information Technology, 13(32), 145–159.
65. Suura, S. R. (2026). Deep Learning-Enabled Cell-Free DNA Analytics for Precision Reproductive and Preventive Healthcare. International Journal of Science, Research and Technology, 9(3), 801-815.
66. Rozman, A. G. (2024). The power of data: Transforming compliance with anti-money laundering measures in domestic and cross-border payments. Journal of Payments Strategy & Systems, 18(3), 253–260.
67. Deloitte. (2025). Agentic AI in banking: Assessing autonomous AI-related risks in financial services. Deloitte Insights.
68. Nandan, B. P. (2026). Cloud-Scale Data Engineering for Real-Time Semiconductor Testing and AI-Powered Chip Diagnostics. International Journal of Computer Technology and Electronics Communication, 9(3), 1058-1070.
69. Aljunaid, S., Alshamrani, A., & Khan, M. (2025). Secure and transparent banking: Explainable AI-driven federated learning model for financial fraud detection. Journal of Risk and Financial Management, 18(4), 179.
70. Mahadevan, S. (2024). Clouding the Future: Innovating Towards Net-Zero Emissions. International Journal of Computing and Engineering, 6(2), 17-23.