ADROID Conference Series: Engineering Sciences and Technology
Volume 2, pp.237-254, 2026

Foundation of Self-Learning AI for Drones

Pragya Saraswat[1], Aditi[2], Satakshi Aishwarya[3], Puneet Garg[4]

[1]SAITM, Gurugram-Delhi (NCR), India

[4]KIET (Deemed to be University), Delhi NCR, Ghaziabad, India

saraswatpragya66@gmail.com, aditi02shukla@gmail.com, satakshi_253173@saitm.ac.in, puneetgarg.research@gmail.com

Abstract: Full autonomy and the end of manual drone operation require a powerful self-training AI spine. Drones piloted with Static Commands would break in unpredictable environments (e.g., windy weather, moving obstacles). Classical self-learning systems establish these principles with strict rules; instead, a drone acquires its abilities through on-policy experience. The drone is a digital pupil that learns from experience. It's mathematically rewarded for getting somewhere and penalised for mistakes, like moving too close to a wall. With thousands of rounds of training under its belt, the AI learns an intricate picture of how to handle flight physics that could not have been hand-coded by a person. Rather than relaying the information to a remote server, the drone performs image processing on the cameras and sensors themselves in real time, with the help of powerful internal chips. These offline computations are complemented by online "Sim-to-Real" training, where the AI is first taught to fly in a safe virtual environment before transfer to physical hardware. That means the drone is already an ace at flying before it ever took off. Also, the foundation must enable lifelong learning. The drone can now adapt to changing environments or mechanical ageing over time, without forgetting what it was taught initially. Developing drones on such intelligent platforms makes them reliable partners for challenging missions, including emergency response and autonomous delivery. This innovation ensures the sky's future will be safe, fluid, and self-determining.

Keywords: Autonomous Drones, Deep Reinforcement Learning, Intelligent Navigation, Sensor Fusion, edge AI, explainable AI, UAV perception.

1. Introduction:

Unmanned aerial systems have been radically transformed over the past 10 years. The once-remotely piloted military reconnaissance platforms have become a panoply of commercially available systems for precision agriculture, infrastructure inspection, search-and-rescue, last-mile logistics, and environmental monitoring. The narrative underlying all of this growth is the gradual transfer of decision-making power away from human operators and into onboard software, a story that ends at what we call full autonomy [1,3]. Such full autonomy requires that a drone sense its environment with a certain level of reliability, plan its trajectories in a dynamic and partially observable environment, and perform actions safely in a real-time bounded environment, while considering the limited energy supply of a battery-powered aircraft. Elegant model-based control principles, however, have difficulty when the surrounding world is unstructured, GPS-denied, or filled with unknown obstacles [73,74].

Nowhere is this truer than in machine learning, and RL in particular: instead of explicitly programming for every contingency, the system learns behavioural policies through interactions with its environment [75,76]. Self-learning AI for drones is not a single algorithm but a composite system [2.3,4]. At the perception layer, deep convolutional and transformer-based models consider the raw sensor streams to produce semantic interpretations. Probabilistic filters at the state-estimation layer fuse diverse sensor measurements into consistent world representations [77,78]. At the control layer, RL agents take world states as inputs and output actions to optimise long-horizon objectives (e.g., mission completion, energy efficiency, collision avoidance). Spanning all layers is the challenge of uncertainty in sensor measurements, in the actions of dynamic agents, and even in the model of the world itself [5].

2. Theoretical Foundations of Self-Learning in Aerial Systems

2.1 Reinforcement Learning Primer

RL represents the problem of sequential decision making through the Markov Decision Process (MDP) framework described by the tuple (S, A, P, R, γ), in which the S represents the state (S1, S2, …, Sn), the A represents the action (A1, A2, …, Am), the P is the transition probability function, the R is the reward function, and $\gamma \in (0,1)$ is the discount factor, which defines the importance of future rewards versus immediate rewards [79,80]. An RL agent is learning a policy $\pi: S \rightarrow A$ that maximises the expected cumulative discounted reward — the return — over episodes while interacting with the environment [6.7].as per fig 1

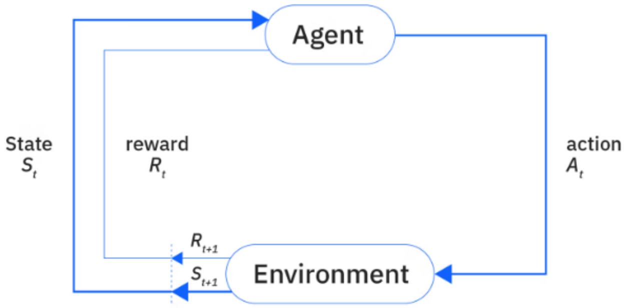

Fig. 1, Reinforcement Learning - Markov decision processes

For UAV, the state space often encodes the vehicle pose (position, orientation, velocity, angular velocity), a local occupancy representation (e.g., a voxel grid generated by depth images), and semantic information relevant to the task (goal position, waypoint progress) [81,82]. The space of actions may be specified at varying degrees of granularity: upon the bottom level, direct commands of motor thrust; at the middle level, desired velocity or acceleration set points that are fed to a classical controller; or at the top level, semantic navigation instructions such as "follow corridor" or "ascend to waypoint" [8,9]. Hierarchical RL schemes take advantage of this property, learning distinct policies at different levels of abstraction and combining them online [83,84].

2.2 Value-Based and Policy-Gradient Methods

Value-based algorithms, including Deep Q-Networks (DQN) and its variants (Double DQN, Prioritised Experience Replay, Duelling DQN), learn an estimate of the action-value function Q(s, a), representing the expected return of executing action a in state s and then following the optimal policy. The recursive structure is provided by the Bellman optimality equation: Q*(s, a) = E[r + γ max_{a'} Q*(s', a') as per figure 2 [10,11].

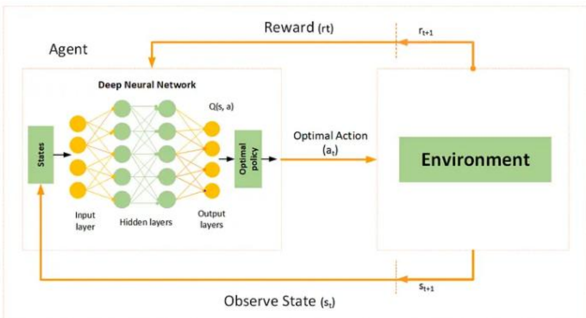

Fig. 2, Structure of DQN

Actor-critic architectures, such as Advantage Actor-Critic (A2C), Asynchronous Advantage Actor-Critic (A3C), and the popular Proximal Policy Optimisation (PPO), decrease variance by learning a value function baseline in parallel with the policy [85,86]. Due to its sample efficiency, robustness against hyperparameter variations, and support for continuous action spaces, PPO is considered the default standard in drone RL studies [12,13].

2.3 Model-Based Reinforcement Learning

An inherent drawback to the model-free RL is that it is sample-inefficient, i.e. it may require millions of interactions with the environment to learn good policies. Model-based reinforcement learning (MBRL) tackles this by learning a dynamics model $f: S \times A \rightarrow S$ and employing it to plan (e.g., through model predictive control), or to generate synthetic rollouts to complement real experience. Approaches like Dreamer, MBPO, and PETS have established new levels of sample efficiency on continuous control benchmarks as per fig 3 [14,15,16].

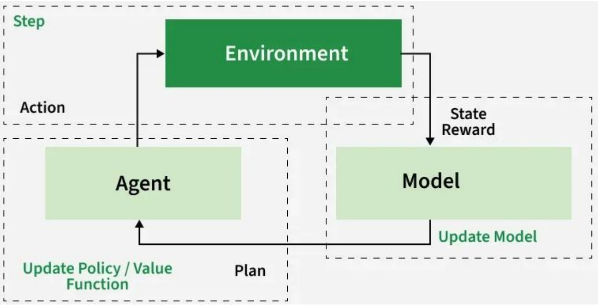

Fig. 3, Model-Based Reinforcement Learning (MBRL)

For drone systems, where collecting data in the real world is costly and potentially destructive, MBRL provides substantial practical benefits [17,18,19].

2.4 Reward Function Design

Since the reward function encodes the objectives of the task, it also defines the types of behaviours that are learned for that task. Mis-specified rewards and the consequent reward hacking, i.e., the agent is finding ways to game the reward function that are contrary to its spirit [21]. For the drones to be autonomous in a sense, the composite reward functions are crafted so as to represent a compromise between the sometimes competing demands of mission priority (distance to goal), safety (cost for near misses, and collisions), energy efficiency (cost for high thruster output), and grace of flight (cost for trajectories that involve more abrupt changes in direction, which tend to impart high inertia to mechanical components) as per fig 4 [22,23].

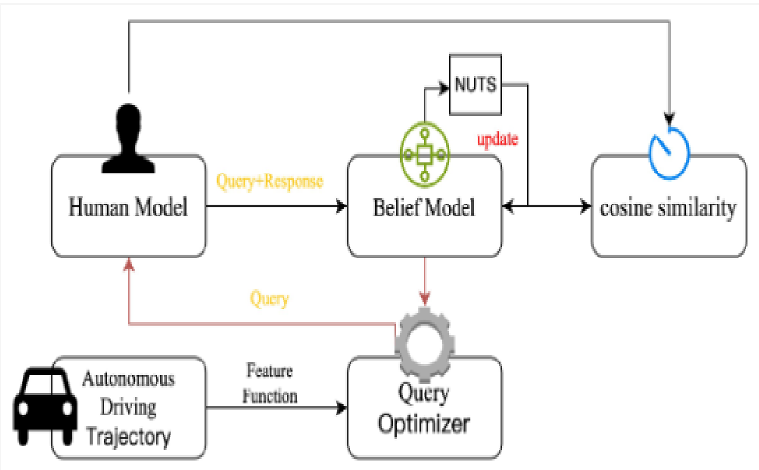

Fig. 4, Reward Function Design

Inverse reinforcement learning (IRL) provides a different approach: instead of manually designing the reward, IRL infers the reward function from expert demonstrations, allowing the transfer of human piloting skills to learning agents [87,88,24].

3. Perception Architecture and Sensor Fusion

3.1 Sensor Modalities

Autonomous drones utilise different types of sensors to form situational awareness. Monocular and stereo RGB cameras capture rich visual texture and colour information with minimal cost and weight [89,90]. Depth cameras (structured-light or time-of-flight) provide a dense range map per-pixel at short distances. LiDAR sensors enable accurate long-distance 3D point clouds, but bring a huge size, weight and power (SWaP) burden [91,92].

Fig. 5, Basics of Sensors in UAVs

Inertial measurement units (IMUs) offer high-rate linear acceleration and angular velocity measurements that are essential for state estimation at rates higher than that of camera frames [93,94]. Barometers and ultrasonic sensors provide altitude information with low processing complexity [25]. Event cameras – neuromorphic sensors that asynchronously report per-pixel brightness changes at a temporal resolution of a few microseconds – are increasingly used for high-speed flight and high dynamic range applications as per fig 5 [26].

3.2 Deep Learning for Visual Perception

Convolutional neural networks (CNNs) have replaced hand-crafted feature extractors for nearly every drone perception problem. For obstacle detection and semantic segmentation, encoder-decoder networks such as U-Net and DeepLab show comparable performance based on dense pixel predictions at real-time speeds when executed on embedded GPU platforms [27,28]. Other object detection models (based on the YOLO family - YOLOv8, YOLOv9, one of the most well-known and used in this field, and DETR - Detection Transformer) are able to achieve state-of-the-art accuracy-latency trade-offs on aerial images [95,96,97]. Depth estimation from monocular imagery, long believed to be ill-posed, has been made feasible by self-supervised learning methods that exploit photometric consistency among video frames, and consist of metric-scale depth maps without the use of LiDAR ground truth [29,30,31].

3.4 Multi-Modal Fusion Strategies

Simultaneous Localisation and Mapping (SLAM) is the spatial enabler for GPS-free autonomous navigation. Visual-inertial odometer (VIO) systems such as VINS-Mono and KIMERA couple camera and IMU measurements tightly via non-linear optimisation and achieve centimetre-level localisation accuracy with real-time frame rates [98,99]. LiDAR-inertial SLAM systems (LIO-SAM, FAST-LIO) extend this approach into 3D point cloud registration and mapping of expansive scenes, forming globally consistent maps of large-scale scenes [100,101]. Learning-based SLAM algorithms, e.g. DeepFactors and iMAP, leverage neural implicit representations to compactly encode maps and enable differentiable map querying as per fig 6. [32,33].

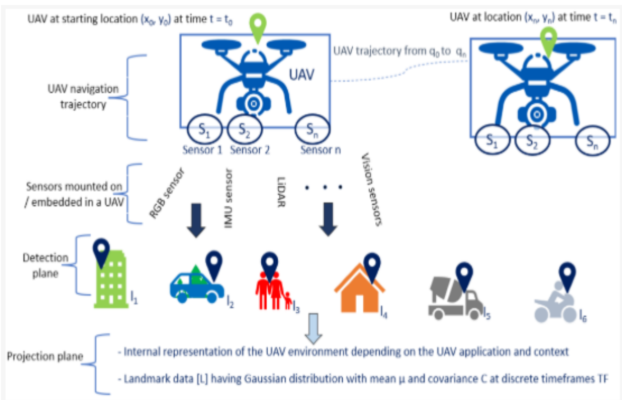

Fig. 6, Multi-Modal Fusion

4. Intelligent Navigation in Complex Environments

4.1 Path Planning Paradigms

Navigation involves global path planning (finding a collision-free path from start to goal in a known map) and local motion planning (generating execution commands to follow the global path in a short time while reacting to unexpected obstacles) [102,103]. Classical global planners such as A*, RRT* and PRM are known to be complete and optimal under idealised assumptions, yet they do not scale well for high-dimensional configuration spaces and have a hard time dealing with dynamic environments [104,105]. Learning-based planners mitigate the above limitations by learning neural networks that output near-optimal paths directly from sensor observations, amortising planning cost across episodes as per fig 7.[34,35].

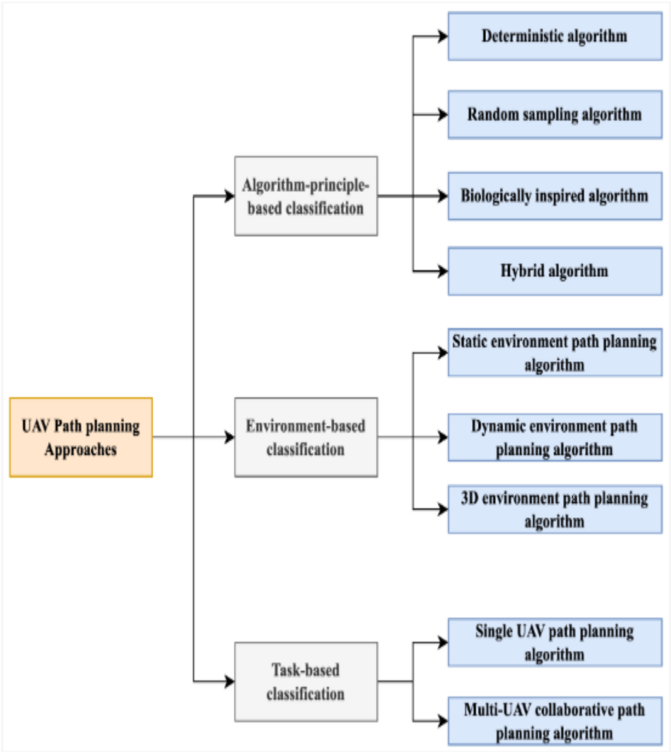

Fig. 7, UAV path planning methods

4.2 End-to-End Navigation with Deep RL

End-to-end RL navigation policies take raw sensor observations as input and directly output flight commands without any explicit mapping [106,107,108]. Pioneering results at ETH Zurich showed that PPO policies learned to navigate forest trails from monocular camera input alone at rates comparable to or faster than expert human pilots, considering both speed and safety [36,37]. Follow-up work generalised this to indoor spaces (using depth cameras) and dynamic urban airspaces (using LiDAR-image fusion)[38,39,40].

4.3 GPS-Denied and Indoor Navigation

Navigation in GPS-denied environments, such as underground or urban canyons, or in warehouse settings, is one of the toughest regimens to work within. Without a reference for its absolute position, the quadrotor must perform state estimation on its own [41,109,110]. Deep RL approaches for this problem commonly store a memory of previous observations through recurrent networks (LSTM, GRU) or transformer decoders, allowing the policy to resolve global pose ambiguity conditioned on the sequence of local observations [42.111,112]. Topological maps (which represent environments as graphs of distinctive places rather than metric coordinates) provide a high-level, robust, and compact representation for large-scale GPS-denied navigation compared with dense occupancy grids as per fig 8[43,113].

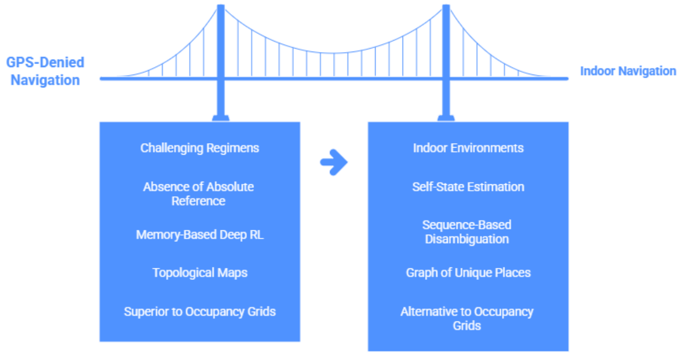

Fig. 8, GPS-Denied and Indoor Navigation

4.4 Dynamic Obstacle Avoidance

Dynamic obstacle avoidance — such as pedestrians, cars, and other UASs— differs fundamentally from static because it requires predicting every agent's future actions [114,115]. Learning-based prediction models identify their obstacles by predicting future states, from simple constant-velocity models to more advanced social force models, and graph neural network-based trajectory predictors provide probabilistic predictions on obstacle dynamics [116,117]. Such predictions may also be built into MPC or risk-aware RL policies (that consider a distribution over future obstacle locations rather than a point estimate) to create proactive avoidance manoeuvres that foresee possible collisions [44,45].

5. Decision-Making under Uncertainty

5.1 Partially Observable Markov Decision Processes

In fact, the drone has access only to noisy and partial observations of the true state of the world — a scenario encapsulated as the Partially Observable Markov Decision Process (POMDP). POMDP planning maintains a belief distribution over the world states, and selects actions so as to maximise the expectation of value within the context of this uncertainty [46,47]. While exact POMDP solutions are computationally intractable for large state spaces, approximate algorithms like POMCP (Monte Carlo tree search over beliefs) and neural belief encoders have been demonstrated to be computationally viable for tasks with high-dimensional observations [118,119,120].

5.2 Multi-Objective Optimisation

In fact, the real drone missions require a trade-off between conflicting goals: minimising mission duration is inconsistent with minimising energy consumption, or with maximising coverage, and minimising risk [50,51]. Multi-objective RL (MORL) extends the scalar reward setup to multiple objectives and learns a Pareto Front of policies that reflect different trade-off solutions [121,122]. Conditioned reward networks allow a policy network to make a smooth decode on the Pareto front (it receives as input a weight vector) and thus let the mission priorities be determined during the mission without retraining [52,123]. Constrained Markov Decision Processes (CMDPs) offer yet another formalism, which tries to optimise the main objective while obeying hard/soft constraints on the secondary objectives, e.g., collision risk or energy budget [124,125].

5.3 Hierarchical and Task-Conditioned Policies

Temporal structure at multiple scales is also present in complex missions such as search and rescue, multi-site delivery, and persistent environmental monitoring [50,51,52]. Hierarchical RL breaks down these types of tasks to a hierarchy of managers and workers: the high-level manager will provide subgoals or subtask identifiers, then low-level workers will take primitive actions to finish them [126,127]. Such decomposition is formally treated within the Options framework in feudal RL (where a manager is also a worker with a higher-level manager) [53,127]. Language-conditioned policies, which take natural language mission specifications as input, are also on the rise: large language models capture the semantics of textual mission descriptions in task embedding conditioning the navigation policy, thus allowing non-expert users to program sophisticated drone behaviours via natural dialogue [54,55].

6. Multi-Agent Systems and Swarm Intelligence

6.1 Cooperative Multi-Agent RL

Fleets of autonomous drones performing collective tasks — such as area coverage, cooperative mapping, and coordinated search— are more likely to be deployed with multi-agent reinforcement learning (MARL) [128,129]. In cooperative MARL, agents receive a single team reward and are required to implicitly coordinate by optimising their joint policy or explicitly through communication [130,131]. This centralised training, decentralised execution (CTDE) paradigm has become the prevailing approach: the training phase allows the agents to have access to each other's observations and states through a centralised critic, but each agent is only able to use its own local observations during execution. QMIX and MADDPG are popular algorithms under this category [56].

6.2 Swarm Communication and Emergence

Human Swarm Park Biologic swarms —such as birds, fish, and insects— display staggering collective intelligence with only simple local interaction rules and no centralised coordination. Today, engineered drone swarms mimic this paradigm more and more – and employ graph neural networks (GNNs) to represent the inter-agent connections [57]. Each drone has access to a local observation of its neighbourhood, and a shared GNN policy predicts over this neighbourhood graph to actions, naturally accounting for changing fleet sizes and topology changes as a result of communication dropouts. Emergent behaviours like formation flight, cooperative load carrying, and distributed search are observed in these learned local policies, in the absence of explicit programming as per fig 9[58,59].

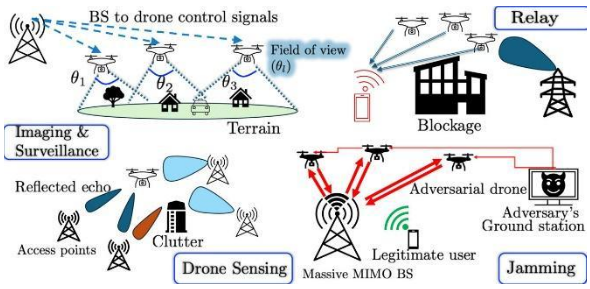

Fig. 9, Interacting with Drone Swarms

6.3 Federated Learning for Collaborative Intelligence

Swarm operation produces huge amounts of scattered data—each drone builds up its own flight experience—but privacy, bandwidth, and regulatory barriers might prevent that data from being consolidated [132,133]. Federated learning (FL) facilitates collaborative enhancement of models by aggregating gradient updates instead of raw data [134,135]. In federated drone systems, a parameter server periodically retrieves model gradient updates from the fleet members, aggregates them (e.g., by FedAvg), and then broadcasts the refined global model [136,137]. Challenges relevant to the drone FL scenario include non-IID data distributions (drones in different environments capture different data), intermittent communication links (drones may be temporarily out of communication range), and requiring Byzantine-robust aggregation techniques to resist adversarial fleet members [60,138].

7. Sim-to-Real Transfer and Curriculum Learning

7.1 The Reality Gap

High-fidelity simulation is crucial for RL policy training: gathering millions of real-world interactions is cost-prohibitive, and early training policies will perform dangerous flights that could damage hardware [139,140]. However, policies trained in sim do not generally work when transferred to real drones due to the reality gap – differences between simulated and real physics, sensor models, motor dynamics and environmental appearance [141]. Closing this gap is a major engineering challenge in applied drone RL [61,62].

7.2 Domain Randomisation and Adaptation

Domain randomisation (DR) tackles the reality gap by training on a large spectrum of simulated worlds — randomising physical parameters (motor thrust coefficients, drag, and mass distribution), sensor noise statistics, lighting conditions, and visual textures [142,143]. The resulting policy would perform well on this range of variation—generalises to the real world, as the real world can be seen as one more sample of the randomisation distribution [144,145]. Adaptive domain randomisation (ADR) builds on this by gradually increasing the randomisation range over training from near-nominal conditions to more challenging conditions [63,64]. Domain adaptation techniques, e.g., CycleGAN−based image translation to map real camera images into the simulation domain (or vice versa), can enhance randomisation by mitigating the visual gap directly without demanding photorealistic rendering [146].

7.3 Curriculum Learning

Training RL policies from scratch in complex environments gives rise to the credit assignment problem, which means that reward signals are too sparse for the agent to infer which behaviours lead to success or failure. Curriculum learning solves this by shaping the training distribution — with simple examples to begin with, and then incrementally adding complexity. For drone navigation, a natural curriculum is to start with wide, obstacle-free corridors, then gradually narrow them and place obstacles as the policy improves. Self-paced curriculum learning approaches, which eliminate this manual design by estimating when an agent is ready to face harder tasks based on performance measures [65, 46]. Hindsight experience replay (HER) is a complementary approach for sparse-reward goal-conditioned tasks: unsuccessful episodes are relabelled offline with the goal that was actually achieved as if it were a positive outcome, significantly increasing the learning signal density [147].

7.4 Simulation Platforms

Several high-fidelity simulators exist and are commonly used for research with drones in RL. Microsoft AirSim simulates multirotor physics and camera sensors in photorealistic Unreal Engine environments. Gazebo (de facto for ROS) has good sensor simulation, and also integrates nicely with your flight controller firmware (PX4, ArduPilot). Flightmare [17], from ETH Zurich, focuses on high-speed rendering, achieving millions of simulated steps per hour by separating the physics and the visual rendering as per fig 10[66].

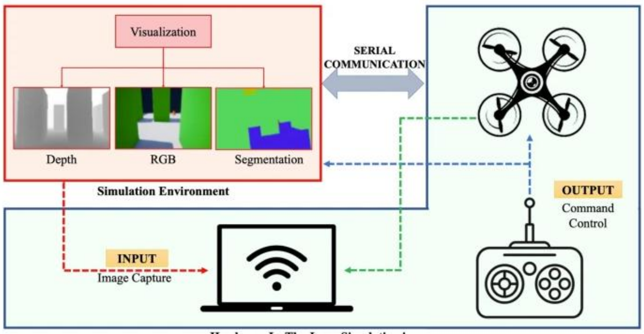

Fig. 10, Drone Simulator

Isaac Gym (NVIDIA) is a GPU-based simulator that exploits GPU parallelism for simulating thousands of physics environments concurrently, resulting in sample throughputs that allow training to be reduced from days to hours. These platforms have evolving support for procedural scenes generation (automatic creation of a variety of training environments at scale) [35,67].

8. Safety, Explainability, and Regulatory Considerations

8.1 Safe Reinforcement Learning

Standard RL methods, which arbitrarily explore the state space and therefore visit many unsafe states, must not be directly applied to hardware, as this would bring unacceptable risks to physical systems in the air. Safe RL methods are often formulated as extensions of the standard MDP framework with the addition of safety constraints or shields [12,55,68]. Constrained Policy Optimisation (CPO), a trust-region based approach to safe RL in the presence of known worst-case cost bounds and unknown entropy bounds, and the Lagrangian relaxation framework that trades off task performance with probabilistic safety constraints satisfaction [148,149]. Safety shields act on RL policy outputs and replace the actions when a classical safety monitor foresees a constraint violation with high probability — leveraging the flexibility of learned policies and the verifiability of formal methods. Control Barrier Functions (CBFs) offer an elegant way to enforce safety that can be certified for ensuring the forward invariance of safe sets of operation [69,70].

8.2 Explainable AI for Aerial Systems

Increasingly, as autonomous drones are being used in safety-critical applications, regulators, operators, and the public are expecting not only high performance but also interpretability — a way to understand why the system made a particular decision. Saliency map methods (GradCAM, Integrated Gradients) highlight which parts of an image have the most influence on the policy action with post-hoc explanations for individual decisions [25,31,56]. Visualising attention weights in transformer-based policies sheds light on what part of the observed scene and task history the agent is attending to when selecting its actions [150]. Concept bottleneck models with human-interpretable intermediate representations, corresponding to concepts, and allow an operator to inspect and intervene on fig 11. [48, 71].

Fig. 11, Autonomous Drones

8.3 Regulatory Landscape

The drone regulation landscape is changing quickly. The U.S. Federal Aviation Administration (FAA) regulates UAV operations under Part 107 rules, and soon other regulations in this framework for BVLOS operation of UAVs will be introduced that will provide the basis for more advanced autonomous applications [130]. The European Union Aviation Safety Agency (EASA) has introduced a risk-based categorisation (Open, Specific, Certified) and associated operational approval conditions [22,41]. For AI-based systems, other challenges come to the fore: much of the regulatory landscape was built around deterministic, human-codified behaviour, and there are few established structures for certifying learned policies with performance guarantees that are probabilistic by nature. RTCA (DO-178C successor) and IEEE emerging standards are also starting to consider AI assurance in safety-critical aerospace systems; the widespread deployment of fully autonomous drones is still dependent on regulatory frameworks aligning with this technology [50,52].

9. Benchmark Evaluations and System Performance

9.1 Benchmark Environments

Several standardised interfaces make it possible to reproduce evaluations of autonomous drone algorithms. The DARPA Subterranean Challenge offered a broad-scale real-world benchmark for GPS-denied navigation in mines, tunnels, and urban underground environments [35,70]. Alpha Pilot and the Autonomous Drone Racing League (ADL) created the high-speed challenge for navigating through obstacle courses, challenging the limits of perception and control algorithms. The MRS Drone Racing and NIST UAV Indoor Navigation datasets also offer rigid evaluation protocols for laboratory work. For multi-agent systems, fleet coordination is tested by the Multi-Robot Exploration (MRE) and the Cooperative Aerial Coverage benchmarks as per fig 12.[59,71].

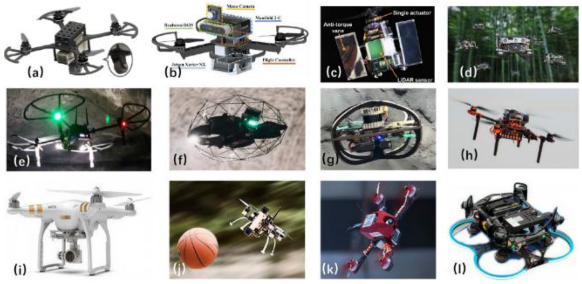

Fig. 12, Multi-Robot Exploration (MRE)

9.2 Key Performance Indicators

Autonomous drone systems evaluation includes more than one aspect. Navigation accuracy is defined in terms of position root mean square error (RMSE) with respect to ground truth trajectories from a motion capture system or differential GPS [116]. Performance of obstacle avoidance, i.e. collision rate (number of collisions per flight hour) and near-miss rate (number of incidents falling within a predefined safety margin), is also predicted. The mission success rate defines the probability of accomplishing the task end-to-end [120]. Efficiency of energy is manifest by mission range (km/Wh), hover endurance, and energy cost of work [22–24]. In addition, the computational cost-related measures, inference latency, operations per second, and power draw of the AI accelerator are also important from a perspective of embedded deployment based on our experience. Sample efficiency in training is reported as performance after a certain number of interactions with the environment or wall-clock training time [31,57].

9.3 Representative Results

Our results show the state-of-the-art performance on self-learning drone systems. In this work, we show in high-speed racing scenarios that RL-trained policies achieve lap times competitive with professional human pilots on the same physical track, which establishes that learned controllers can achieve human or even super-human level performances in speed-critical tasks [134]. In GPS-denied exploration, a team of collaborative drone fleets, controlled through GNN-based coordination policies, explored underground environments 40% faster than single-drone baselines while retaining the equivalent mapping accuracy [135]. Sim-to-real transfer experiments showed policies trained with domain randomisation on 50,000 simulated environments transferred to real drones with under 15% performance degradation — a drastic improvement over policies trained on a solitary fixed simulation, which degraded by over 60% [30,66].

10. OPEN RESEARCH CHALLENGES

10.1 Long-Horizon Mission Planning

Current RL algorithms perform well on short-horizon tasks but have difficulty planning and executing missions over the course of hours or days that involve dozens of subtasks and require the integration of high-level semantic knowledge with low-level motor control [136]. Closing this gap will require progress in hierarchical RL, memory-augmented neural architectures, and world models that can predict the evolution of the environment over long time horizons. Foundation models pre-trained on internet-scale data — such as large language models and vision-language models — may contain the rich world knowledge needed to enable long-horizon planning, but how they can be integrated with real-time flight control is largely an open question [64].

10.2 Adversarial Robustness

Deep learning elements of drone perception informatics are susceptible to adversarial perturbations — maliciously modified inputs that a human would not notice but that cause a system to malfunction. In relation to drones, adversarial attacks could be physically implemented, for instance, by using camouflage on objects or lasers to interfere with camera sensors [52,46]. Certified robustness techniques such as randomised smoothing and interval bound propagation come with theoretical robustness guarantees against bounded perturbations, but usually at the expense of accuracy. Achieving engineered solutions for which perception systems are certifiably robust and performant under natural distribution shift as well as intentional adversarial perturbations is a major unsolved problem for security-sensitive applications [29,70].

10.3 Ethical Considerations and Dual Use

The emergence of fully autonomous drones with far-reaching AI capabilities raises deep ethical concerns. The same tools that afford precision delivery of humanitarian assistance can be turned into those for targeted surveillance, or lethal engagement, in an automated fashion [125]. Payload-agnostic AI capabilities — perception, navigation, decision-making — are dual-use to their core, and the international community does not have a clearly established norm or agreement on how these should be developed and applied [127]. Autonomous drone AI research organisations and companies must proactively engage with ethicists, policymakers, and civil society to outline responsible research and deployment practices, including impact assessments, use-case restrictions, and transparent disclosure of capabilities and limitations as per fig 13.[47,63].

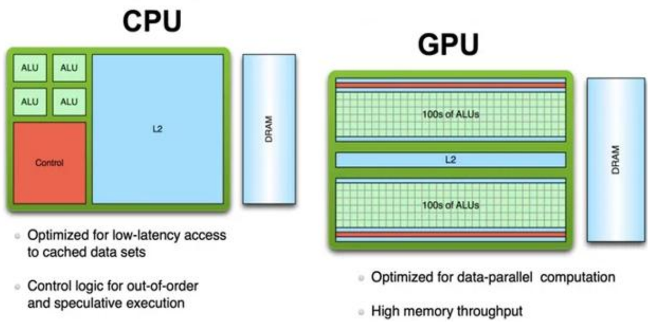

Fig. 13, GPU- based Drones

10.4 Energy Efficiency and Extended Endurance

Battery-powered drones are currently severely limited in endurance: typical commercial multirotor platforms only achieve 20–40 minutes of flight per charge. The computational resources required for on-board AI also drain batteries, introducing a fundamental trade-off between intelligence and endurance [37,40]. Neuromorphic computing architectures, which perform computation in an asynchronous spiking manner mimicking biological neural systems, can lead to orders-of-magnitude enhancement in energy efficiency compared to conventional GPU-based inference methods. Pruning, quantisation and knowledge distillation reduce model size, allowing inference on-device for powerful AI on microcontroller-class hardware[71,72].

11. CONCLUSION

Self-learning AI is the biggest breakthrough in autonomous drone technology in the last decade. The combination of deep reinforcement learning, multi-modal perception, and multi-agent coordination has resulted in systems that enable unprecedented levels of autonomy, operating in GPS-denied areas, dynamically avoiding obstacles, and performing complex multi-stage missions that five years ago would have been considered aspirational. This chapter has followed the theoretical and algorithmic machinery underlying these advances - from the formal MDP framework and its extensions to partial observability and safety constraints, over the perception architectures that ground policy learning in rich sensory reality, all the way to the simulation platforms and curriculum learning techniques that enable scaling of safe training. We have explored how swarm intelligence and federated learning extend these capabilities to collaborative multi-drone systems, and how explainability and formal safety methods are opening the door to meeting the rigours of safety-critical certification. Several themes were noted as being consistently significant in the literature. The quality and variety of simulation environments are the most important factor influencing real-world performance — investment in simulation fidelity and procedural scene generation results in compounding returns in downstream policy quality. Second, uncertainty — in sensing, modelling, and the behaviour of other agents — must be treated as a first-class concern rather than an afterthought in the algorithmic design of perception and decision systems and in the formal safety analysis required for deployment approval.

Third, the human interface is still developing in relation to the autonomy of the systems themselves — the joystick pilot is becoming a mission commander, requiring new paradigms for operator engagement, situation awareness, and the management of authority. Looking ahead, combining foundation models -- large-scale vision-language and world models pre-trained on internet-scale data -- with real-time flight control may well be the greatest short-term research opportunity. These models provide rich priors for scene understanding, goal inference from natural language, and generalisation to novel environments, which could greatly extend the operational envelope of autonomous drone systems. Neuromorphic computing and on-device learning, on the other hand, hold the promise of overcoming the energy limits that currently restrict the amount of intelligence that can be deployed at the edge.

References:

[1] Saraswat, P., Garg, P., & Siddiqui, Z. (2025). AI & the Indian Stock Market: A Review of Applications in Investment Decision. In the *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. https://doi.org/10.63169/GCARED2025.p10

[2] Aditi, Garg, P., & Roy, B. (2025). A System of Computer Network: Based On Artificial Intelligence. In the *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. https://doi.org/10.63169/GCARED2025.p24

[3] Vignesh, R., & Samydurai, A. (2017). Security in the Internet of Things (IoT): challenges and countermeasures. IJEDR, 5(1).

[4] Koblitz, N. (1987). Elliptic curve cryptosystems. Mathematics of Computation, 48, 203–209.

[5] Lee, J.-Y., Lin, W.-C., & Huang, Y.-H. (2014). A lightweight authentication protocol for the Internet of Things. In Proceedings of ISNE (pp. 1–2).

[6] Pustokhina, I. V., Pustokhin, D. A., Lydia, E. L., Garg, P., Kadian, A., & Shankar, K. (2021). Hyperparameter search-based convolutional neural network with Bi-LSTM for an intrusion detection system in a multimedia big data environment. Multimedia Tools and Applications, 1–18.

[7] Khanna, A., Rani, P., Garg, P., Singh, P. K., & Khamparia, A. (2021). An enhanced crow search-inspired feature selection technique for intrusion detection in a wireless network system. Wireless Personal Communications, 1–18.

[8] Garg, P., Dixit, A., & Sethi, P. (2020). Impact of node density on the QoS parameters of routing protocols in opportunistic networks for smart spaces. Mobile Information Systems.

[9] Upadhyay, D., Garg, P., Aldossary, S. M., Shafi, J., & Kumar, S. (2023). A linear quadratic regression-based synchronised health monitoring system (SHMS) for IoT applications. Electronics, 12(2), 309.

[10] Saini, P., Nagpal, B., Garg, P., & Kumar, S. (2023). CNN-BI-LSTM-CYP: A deep learning approach for sugarcane yield prediction: Sustainable Energy Technologies and Assessments, 57, 103263.

[11] Saini, P., Nagpal, B., Garg, P., & Kumar, S. (2023). Evaluation of remote sensing and meteorological parameters for yield prediction of sugarcane crop—Brazilian Archives of Biology and Technology, 66, e23220781.

[12] Abomhara, M., & Køien, G. M. (2014). Security and privacy in the Internet of Things: Current status and open issues. In Proceedings of PRISMS (pp. 1–8). IEEE.

[13] Chen, S., Xu, H., Liu, D., Hu, B., & Wang, H. (2014). A vision of IoT: Applications, challenges, and opportunities with China. IEEE Internet of Things Journal, 1(4), 349–359.

[14] Zanella, A., Bui, N., Castellani, A., Vangelista, L., & Zorzi, M. (2014). Internet of Things for smart cities. IEEE Internet of Things Journal, 1(1), 22–32.

[15] Miorandi, D., Sicari, S., De Pellegrini, F., & Chlamtac, I. (2012). Internet of things: Vision, applications and research challenges. Ad Hoc Networks, 10(7), 1497–1516.

[16] Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. Computer Networks, 54(15), 2787–2805.

[17] Amendola, S., Lodato, R., Manzari, S., Occhiuzzi, C., & Marrocco, G. (2014). RFID technology for IoT-based personal healthcare in smart spaces. IEEE Internet of Things Journal, 1(2), 144–152.

[18] Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. IEEE Communications Surveys & Tutorials, 17(4), 2347–2376.

[19] Sharma, K. K., Verma, P. K., & Garg, P. (2024). IoT-enabled energy management systems for sustainable energy storage: Design, optimisation, and future directions. Frontiers in Health Informatics, 13(8).
[20] Garg, P., Dixit, A., & Sethi, P. (2019). Wireless sensor networks: An insight review. International Journal of Advanced Science and Technology, 28(15), 612–627.
[21] Sharma, N., & Garg, P. (2022). Ant colony-based optimisation model for QoS-based task scheduling in cloud computing environment—measurement: Sensors, 100531.
[22] Garg, P., Dixit, A., & Sethi, P. (2022). ML-fresh: Novel routing protocol in opportunistic networks using machine learning. Computer Systems Science & Engineering, Tech Science Press.
[23] Beniwal, S., Garg, P., Rajpal, R., Sharma, N., & Mittal, H. K. (2025). Fusion of opportunistic networks with machine learning: Present and future. Metallurgical and Materials Engineering, 31(1), 311–324.
[24] Garg, P. (2025). Explainable AI & model interpretability in healthcare: Challenges & future directions. EKSPLORIUM, 46(1), 104–133.
[25] Rani, P. (2025). From data to diagnosis: Unleashing AI and 6G in modern medicine. EKSPLORIUM, 46(1), 69–103.
[26] Garg, P., Sharma, N., & Shukla, B. (2023). Predicting the risk of cardiovascular diseases using machine learning techniques. International Journal of Intelligent Systems and Applications in Engineering, 11(2s), 165–173.
[27] Patil, S. C., Mane, D. A., Singh, M., Garg, P., Desai, A. B., & Rawat, D. (2024). Parkinson's disease progression prediction using grey wolf optimiser-based feature selection. International Journal of Intelligent Systems and Applications in Engineering, 12(3s), 441–451.
[28] Rai, S., Choubey, V., Suryansh, & Garg, P. (2022). A systematic review of encryption and keylogging for computer system security. In 2022, Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 157–163). IEEE.
[29] Raj, G., Verma, A., Dalal, P., Shukla, A. K., & Garg, P. (2023). Performance comparison of several LPWAN technologies for an energy-constrained IoT network. International Journal of Intelligent Systems and Applications in Engineering, 11(1s), 150–158.
[30] Beniwal, S., Saini, U., Garg, P., & Joon, R. K. (2021). Improving performance during camera surveillance by integration of edge detection in an IoT system. International Journal of E-Health and Medical Communications, 12(5), 84–96.
[31] Kumar, P., Kumar, R., & Garg, P. (2020). Hybrid crowd cloud routing protocol for wireless sensor networks. International Journal of Advanced Science and Technology, 29, 766–775.
[32] Garg, P., Dixit, A., Sethi, P., & Pinheiro, P. R. (2020). Impact of node density on the QoS parameters of routing protocols in opportunistic networks. Mobile Information Systems.
[33] Sharma, R., Gupta, S., & Garg, P. (2022). Model for predicting cardiac health using a deep learning classifier. In 2022, Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 25–30). IEEE.
[34] Soni, E., Nagpal, A., Garg, P., & Pinheiro, P. R. (2022). Assessment of compressed and decompressed ECG databases for telecardiology applying a convolutional neural network. Electronics, 11(17), 2708.
[35] Yadav, P. S., Khan, S., Singh, Y. V., Garg, P., & Singh, R. S. (2022). A lightweight deep learning-based approach for jazz music generation in MIDI format. Computational Intelligence and Neuroscience.
[36] Garg, P., Dixit, A., Sethi, P., & Pruthi, J. (2023). Strengthening smart cities with opportunistic networks: An insight. In the 2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech) (pp. 700–707). IEEE.
[37] Rana, S., Chaudhary, R., Gupta, M., & Garg, P. (2023). Exploring different techniques for emotion detection through face recognition. In ICACCTech 2023 (pp. 779–786). IEEE.
[38] Mittal, K., Srivastava, K., Gupta, M., & Garg, P. (2023). Exploration of different techniques for heart disease prediction. In ICACCTech 2023 (pp. 758–764). IEEE.
[39] Gautam, V. K., Gupta, S., & Garg, P. (2024). Automatic irrigation system using IoT. In 2024 International Conference on Automation and Computation (AUTOCOM) (pp. 100–103). IEEE.
[40] Ramasamy, L. K., et al. (2024). Forecast of students' mental health combining an artificial intelligence technique and a fuzzy inference system. In AUTOCOM 2024 (pp. 85–90). IEEE.
[41] Rajput, R., Sukumar, V., Patnaik, P., Garg, P., & Ranjan, M. (2024). The cognitive analysis for an approach to neuroscience. In AUTOCOM 2024 (pp. 524–528). IEEE.
[42] Gudur, A., Pati, P., Garg, P., & Sharma, N. (2024). Radiomics feature selection for lung cancer subtyping and prognosis prediction. International Journal of Intelligent Systems and

Applications in Engineering, 12(3s), 553–565.
[43] Dixit, A., Sethi, P., Garg, P., & Pruthi, J. (2024). CNN-based lip-reading system for visual input: A review. AIP Conference Proceedings, 3121(1). AIP Publishing.
[44] Bose, D., Arora, B., Srivastava, A. K., & Garg, P. (2024). A computer vision-based framework for posture analysis and performance prediction in athletes. In IC3SE 2024 (pp. 942–947). IEEE.
[45] Singh, M., Garg, P., Srivastava, S., & Saggu, A. K. (2024). Revolutionising arrhythmia classification: Unleashing the power of machine learning and data amplification. In CCICT 2024 (pp. 516–522). IEEE.
[46] Kumar, R., Das, R., Garg, P., & Pandita, N. (2024). Duplicate node detection method for wireless sensors. In CCICT 2024 (pp. 512–515). IEEE.
[47] Bhardwaj, H., Das, R., Garg, P., & Kumar, R. (2024). Handwritten text recognition using deep learning. In CCICT 2024 (pp. 506–511). IEEE.
[48] Gill, A., Jain, D., Sharma, J., Kumar, A., & Garg, P. (2024). Deep learning approach for facial identification for online transactions. In INNOCOMP 2024 (pp. 715–722). IEEE.
[49] Mittal, H. K., Dalal, P., Garg, P., & Joon, R. (2024). Forecasting pollution trends: Comparing linear, logistic regression, and neural networks. In INNOCOMP 2024 (pp. 411–419). IEEE.
[50] Malik, T., Nandal, V., & Garg, P. (2024). Deep learning-based classification of diabetic retinopathy: Leveraging the power of VGG-19. In INNOCOMP 2024 (pp. 645–651). IEEE.
[51] Srivastava, A. K., Verma, I., & Garg, P. (2024). Improvements in recommendation systems using graph neural networks. In INNOCOMP 2024 (pp. 668–672). IEEE.
[52] Aggarwal, A., Jain, D., Gupta, A., & Garg, P. (2024). Analysis and prediction of customer churn and retention rates in the telecom industry. In INNOCOMP 2024 (pp. 723–727). IEEE.
[53] Mittal, H. K., Arsalan, M., & Garg, P. (2024). A novel deep learning model for effective story point estimation in agile software development. In INNOCOMP 2024 (pp. 404–410). IEEE.
[54] Shukla, S. M., Magoo, C., & Garg, P. (2024). Comparing fine-tuned LMs for detecting LLM-generated text. In DELCON 2024 (pp. 1–8). IEEE.
[55] Kumar, B., Iqbal, M., Parmer, R., Garg, P., Rani, S., & Agrawal, A. (2025). The role of AI in optimising healthcare appointment scheduling. In ICDT 2025 (pp. 881–887). IEEE.
[56] Kumar, B., Garg, V., Ahmed, K., Garg, P., Choudhary, S., & Baniya, P. (2025). Enhancing healthcare with blockchain: Innovations in data privacy, security, and interoperability. In ICDT 2025 (pp. 932–938). IEEE.
[57] Raj, V., Prakash, B. K., Kumar, A., & Garg, P. (2024). Optimise the time a Mercedes-Benz spends on the test bench using stacking ensemble learning. In ICPIDS 2024 (pp. 445–450). IEEE.
[58] Kaushik, N., Kumar, H., Raj, V., & Garg, P. (2024). Proactive fault prediction in microservices applications using trace logs and monitoring metrics. In ICPIDS 2024 (pp. 410–415). IEEE.
[59] Kumar, A. A., Sri, C. V., Bohara, K. S. K., Setia, S., & Garg, P. (2024). Capnivesh: Financing platform for startups. In ICPIDS 2024 (pp. 261–265). IEEE.
[60] Bhandari, P., Setia, S., Kumar, K., & Garg, P. (2024). Optimising cross-platform development with CI/CD and containerization: A review. In ICPIDS 2024 (pp. 175–180). IEEE.
[61] Chaudhary, A., & Garg, P. (2014). Detecting and diagnosing a disease using a patient monitoring system. International Journal of Mechanical Engineering and Information Technology, 2(6), 493–499.
[62] Malik, K., Raheja, N., & Garg, P. (2011). Enhanced FP-growth algorithm. International Journal of Computational Engineering and Management, 12, 54–56.
[63] Garg, P., Dixit, A., & Sethi, P. (2021). Link prediction techniques for opportunistic networks using machine learning, in Proceedings of ICICC.
[64] Garg, P., Dixit, A., & Sethi, P. (2021). Opportunistic networks: Protocols, applications & simulation trends. In Proceedings of ICICC.
[65] Garg, P., Dixit, A., & Sethi, P. (2021). Performance comparison of the fresh and spray-and-wait protocols using a single simulator. IT in Industry, 9(2).
[66] Malik, M., Singh, Y., Garg, P., & Gupta, S. (2020). Deep learning in the healthcare system. International Journal of Grid and Distributed Computing, 13(2), 469–468.
[67] Gupta, M., Garg, P., Gupta, S., & Joon, R. (2020). A novel approach for malicious node detection in cluster-head gateway switching routing in mobile ad hoc networks. International Journal of Future Generation Communication and Networking, 13(4), 99–111.
[68] Gupta, A., Garg, P., & Sonal, Y. S. (2020). Edge detection-based 3D biometric system for the security of web-based payment and task management applications. International Journal of Grid and Distributed Computing, 13(1), 2064–2076.

[69] Garg, P., & Raman, P. K. (2011). Broadcasting protocol & routing characteristics with wireless ad-hoc networks. International Journal of Computer Engineering & Management, 12(1), 36–40.
[70] Garg, P., Arora, N., & Malik, T. (2011). Capacity improvement of WiMAX in the presence of different codes. IJCEM, 12.
[71] Garg, P., Saroha, K., & Lochab, R. (2011). Review of wireless sensor networks: architecture and applications. IJCSMS, 11(01).
[72] Dixit, A., Sethi, P., & Garg, P. (2022). Rakshak: A child identification software for recognising missing children using machine learning-based speech clarification. International Journal of Knowledge-Based Organisations, 12(3), 1–15.
[73] Saraswat, P., & Garg, P. (2026). Soft Computing In AI Agents.
[74] Saraswat, P., & Garg, P. (2026). Water Quality Prediction Using IOT Sensors and Deep Networks.
[75] Arsalan, M., Ahmed, K., & Garg, P. (2026). Machine learning for Anomaly detection in sensor networks. *Available at SSRN 6441518*.
[76] Kumari, M., & Garg, P. (2026). Hybrid Cloud Infrastructure: Models, Benefits, Security, and Challenges. *Benefits, Security, and Challenges (March 19, 2026)*.
[77] Singh, H., & Garg, P. (2026). Demystifying Artificial Distributed Intelligence (ADI). *Available at SSRN 6442698*.
[78] Saraswat, P., & Garg, P. (2026). Human AI Collaboration: The Future of Clinical Decision Making.
[79] Saraswat, P., & Garg, P. (2026). Breaking Data Boundaries: Federated Learning in Digital Healthcare.
[80] Singh, A., Parmar, R., Bhardwaj, P., Sharma, V., & Garg, P. (2026). Fusion of Aerial Networks with Advanced Computing Paradigms. *Edge Computing and Aerial Platforms*, 355-367.
[81] Kumari, M., Baranwal, A., Sonal, & Garg, P. (2026). Application of Aerial Edge Computing in Disaster Management. *Edge Computing and Aerial Platforms*, 103-122.
[82] Aditi, Saraswat, P., Sharma, V., & Garg, P. (2026). Advances in Aerial Platforms and Edge Computing. *Edge Computing and Aerial Platforms*, 123-143.
[83] Garg, P. (2026). Zero-Trust Security Enforcement through AI-Powered Anomaly Detection in Cloud Systems. *Journal of Artificial Intelligence in Governance and Public Policy (JAIGPP)*, *1*(1), 1-8.
[84] Singh, A. P., Sharma, A., & Garg, P. (2026, January). AI-Powered Adaptive Mock Interview Generation System. In *2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC)* (pp. 1421-1426). IEEE.
[85] Raghav, A., Mishra, A., & Garg, P. (2026, January). Enhancing Healthcare Access: An AI-driven Chatbot for Doctor Appointment Management. In *2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC)* (pp. 910-914). IEEE.
[86] Kumar, B., Chauhan, D., Singh, H., Verma, H., Sahu, K., & Garg, P. (2026, January). Decoding Linear A with Artificial Intelligence: A Comprehensive Machine Learning and NLP Framework. In *2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC)* (pp. 1-7). IEEE.
[87] Gupta, V., Chakravarti, L., Akhtar, M. M., Maheshwari, P., Garg, P., & Tiwari, D. (2026, January). A Sentence-Level Risk Estimator for Identifying Hallucinations in Generative AI. In *2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC)* (pp. 1619-1626). IEEE.
[88] Patnaik, P., & Garg, P. (2026). *Principles of Artificial Intelligence: Cognitive Architectures, Large Language Models, and Computational Limits*. Deep Science Publishing.
[89] Singh, K., & Garg, P. (2026). *Trustworthy Deep Learning: Robustness, Uncertainty Quantification, and Adversarial Resilience*. Deep Science Publishing.
[90] Garg, P. (2026). Survey of Load Balancing Strategies in Fog-Cloud Architectures for IoT Integration. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, *9*(2), 595-604.
[91] Tiwari, D., Bhati, B. S., Garg, P., & Babbar, N. (2026). A novel lexicon dictionary and CNN-LSTM employed hybrid approach for sentiment detection of COVID-19 vaccines. *Scientific Reports*.
[92] Sethi, P., Kumar, B., Garg, P., & Yadav, P. (2026, February). Graph Neural Networks for Causal Inference in Climate Science: A Novel Approach to Modelling Complex Interactions. In *2026 International Conference on Intelligent Computing and Automation for Sustainable Solutions (ICASS)* (pp. 1-7). IEEE.

[93] Kumar, B., Garg, P., Garg, A., Nanwal, J., & Yadav, P. (2026, February). An AI-Centric Multimodal Framework for Cognitive Productivity and Digital Wellbeing. In *2026 International Conference on Intelligent Computing and Automation for Sustainable Solutions (ICASS)* (pp. 1-7). IEEE.

[94] Garg, P., Dixit, A., & Sethi, P. (2022). Ml-fresh: novel routing protocol in opportunistic networks using machine learning. *Computer Systems Science & Engineering, Forthcoming*. Tech Science Press.

[95] Yadav, P. S., Khan, S., Singh, Y. V., Garg, P., & Singh, R. S. (2022). A Lightweight Deep Learning-Based Approach for Jazz Music Generation in MIDI Format. *Computational Intelligence and Neuroscience*, *2022*.

[96] Soni, E., Nagpal, A., Garg, P., & Pinheiro, P. R. (2022). Assessment of Compressed and Decompressed ECG Databases for Telecardiology Applying a Convolution Neural Network. *Electronics*, *11*(17), 2708.

[97] Pustokhina, I. V., Pustokhin, D. A., Lydia, E. L., Garg, P., Kadian, A., & Shankar, K. (2021). Hyperparameter search-based convolutional neural network with Bi-LSTM for an intrusion detection system in a multimedia big data environment. *Multimedia Tools and Applications*, 1-18.

[98] Khanna, A., Rani, P., Garg, P., Singh, P. K., & Khamparia, A. (2021). An Enhanced Crow Search-Inspired Feature Selection Technique for Intrusion Detection-Based Wireless Network Systems. *Wireless Personal Communications*, 1-18.

[99] Garg, P., Dixit, A., Sethi, P., & Pinheiro, P. R. (2020). Impact of node density on the QoS parameters of routing protocols in opportunistic networks for smart spaces. *Mobile Information Systems*, *2020*.

[100] Upadhyay, D., Garg, P., Aldossary, S. M., Shafi, J., & Kumar, S. (2023). A Linear Quadratic Regression-Based Synchronised Health Monitoring System (SHMS) for IoT Applications. Electronics, 12(2), 309.

[101] Saini, P., Nagpal, B., Garg, P., & Kumar, S. (2023). CNN-BI-LSTM-CYP: A deep learning approach for sugarcane yield prediction: Sustainable *Energy Technologies and Assessments*, *57*, 103263.

[102] Saini, P., Nagpal, B., Garg, P., & Kumar, S. (2023). Evaluation of Remote Sensing and Meteorological Parameters for Yield Prediction of Sugarcane (Saccharum officinarum L.) Crop. Brazilian Archives of Biology and Technology, 66, e23220781.

[103] Beniwal, S., Saini, U., Garg, P., & Joon, R. K. (2021). Improving performance during camera surveillance by integration of edge detection in an IoT system. *International Journal of E-Health and Medical Communications (IJEHMC)*, *12*(5), 84-96.

[104] Garg, P., Dixit, A., & Sethi, P. (2019). Wireless sensor networks: an insight review. *International Journal of Advanced Science and Technology*, *28*(15), 612-627.

[105] Sharma, N., & Garg, P. (2022). Ant colony-based optimisation model for QoS-Based task scheduling in cloud computing environment—measurement*: Sensors*, 100531.

[106] Kumar, P., Kumar, R., & Garg, P. (2020). Hybrid Crowd Cloud Routing Protocol For Wireless Sensor Networks. *International Journal of Advanced Science and Technology*, *29*, 766-775.

[107] Raj, G., Verma, A., Dalal, P., Shukla, A. K., & Garg, P. (2023). Performance Comparison of Several LPWAN Technologies for Energy-Constrained IoT Network. *International Journal of Intelligent Systems and Applications in Engineering*, *11*(1s), 150-158.

[108] Garg, P., Sharma, N., & Shukla, B. (2023). Predicting the Risk of Cardiovascular Diseases using Machine Learning Techniques. *International Journal of Intelligent Systems and Applications in Engineering*, *11*(2s), 165-173.

[109] Patil, S. C., Mane, D. A., Singh, M., Garg, P., Desai, A. B., & Rawat, D. (2024). Parkinson's Disease Progression Prediction Using Longitudinal Imaging Data and Grey Wolf Optimiser-Based Feature Selection. *International Journal of Intelligent Systems and Applications in Engineering*, *12*(3s), 441-451.

[110] Gudur, A., Pati, P., Garg, P., & Sharma, N. (2024). Radiomics Feature Selection for Lung Cancer Subtyping and Prognosis Prediction: A Comparative Study of Ant Colony Optimisation and Simulated Annealing. *International Journal of Intelligent Systems and Applications in Engineering*, *12*(3s), 553-565.

[111] Khan, A. (2024). Optimisation Methods Based on Soft Computing for Improving Power System Stability. *J. Electrical Systems*, *20*(6s), 1051-1058.

[112] Sharma, K. K., Verma, P. K., & Garg, P. (2024). IoT-Enabled Energy Management Systems For Sustainable Energy Storage: Design, Optimisation, And Future Directions. *Frontiers in Health Informatics*, *13*(8).

[113] Gupta, S., Yadav, N., Singh, K., & Garg, P. (2025). APPLICATIONS OF SIMULATIONS AND QUEUING THEORY IN A SUPERMARKET *Reliability: Theory & Applications*, *20*(1 (82)), 135-140.

[114] Beniwal, S., Garg, P., Rajpal, R., Sharma, N., & Mittal, H. K. (2025). Fusion of Opportunistic Networks with Machine Learning: Present and Future. *Metallurgical and Materials Engineering*, *31*(1), 311-324.

[115] Garg, P. (2025). Explainable AI & Model Interpretability in Healthcare: Challenges & Future Directions. *EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR*, *46*(1), 104-133.

[116] Rani, P. (2025). From Data to Diagnosis: Unleashing AI and 6G in Modern Medicine. *EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR*, *46*(1), 69-103.

[117] Dixit, A., Garg, P., Sethi, P., & Singh, Y. (2020, April). TVCCCS: Television Viewer's Channel Cost Calculation System on Per-Second Usage. In *IOP Conference Series: Materials Science and Engineering* (Vol. 804, No. 1, p. 012046). IOP Publishing.

[118] Sethi, P., Garg, P., Dixit, A., & Singh, Y. (2020, April). Smart number cruncher–a voice-based calculator. In *IOP Conference Series: Materials Science and Engineering* (Vol. 804, No. 1, p. 012041). IOP Publishing.

[119] S. Rai, V. Choubey, Suryansh and P. Garg, "A Systematic Review of Encryption and Keylogging for Computer System Security," *2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT)*, 2022, pp. 157-163, doi: 10.1109/CCiCT56684.2022.00039.

[120] L. Mohanty, L. Saraswat, P. Garg and S. Lamba, "Recommender Systems in E-Commerce," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 114-119, doi: 10.1109/CCiCT56684.2022.00032.

[121] C. Maggo and P. Garg, "From linguistic features to their extractions: Understanding the semantics of a concept," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 427-431, doi: 10.1109/CCiCT56684.2022.00082.

[122] N. Puri, P. Saggar, A. Kaur and P. Garg, "Application of ensemble Machine Learning models for phishing detection on web networks," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 296-303, doi: 10.1109/CCiCT56684.2022.00062.

[123] R. Sharma, S. Gupta and P. Garg, "Model for Predicting Cardiac Health using Deep Learning Classifier," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 25-30, doi: 10.1109/CCiCT56684.2022.00017.

[124] Varshney, S. Lamba and P. Garg, "A Comprehensive Survey on Event Analysis Using Deep Learning," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 146-150, doi: 10.1109/CCiCT56684.2022.00037.

[125] Dixit, A., Sethi, P., Garg, P., & Pruthi, J. (2022, December). Speech Difficulties and Clarification: A Systematic Review. In *2022, the 11th International Conference on System Modelling & Advancement in Research Trends (SMART)* (pp. 52-56). IEEE.

[126] Garg, P., Dixit, A., Sethi, P., & Pruthi, J. (2023, December). Strengthening Smart City with Opportunistic Networks: An Insight. In the *2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech)* (pp. 700-707). IEEE.

[127] Rana, S., Chaudhary, R., Gupta, M., & Garg, P. (2023, December). Exploring Different Techniques for Emotion Detection Through Face Recognition. In *2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech)* (pp. 779-786). IEEE.

[128] Mittal, K., Srivastava, K., Gupta, M., & Garg, P. (2023, December). Exploration of Different Techniques on Heart Disease Prediction. In *2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech)* (pp. 758-764). IEEE.

[129] Gautam, V. K., Gupta, S., & Garg, P. (2024, March). Automatic Irrigation System using IoT. In *2024 International Conference on Automation and Computation (AUTOCOM)* (pp. 100-103). IEEE.

[130] Ramasamy, L. K., Khan, F., Joghee, S., Dempere, J., & Garg, P. (2024, March). Forecast of Students' Mental Health Combining an Artificial Intelligence Technique and Fuzzy Inference

System. In *2024 International Conference on Automation and Computation (AUTOCOM)* (pp. 85-90). IEEE.

[131] Rajput, R., Sukumar, V., Patnaik, P., Garg, P., & Ranjan, M. (2024, March). The Cognitive Analysis for an Approach to Neuroscience. In *2024 International Conference on Automation and Computation (AUTOCOM)* (pp. 524-528). IEEE.

[132] Dixit, A., Sethi, P., Garg, P., Pruthi, J., & Chauhan, R. (2024, July). CNN-based lip-reading system for visual input: A review. In *AIP Conference Proceedings* (Vol. 3121, No. 1). AIP Publishing.

[133] Bose, D., Arora, B., Srivastava, A. K., & Garg, P. (2024, May). A Computer Vision-Based Framework for Posture Analysis and Performance Prediction in Athletes. In *2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE)* (pp. 942-947). IEEE.

[134] Singh, M., Garg, P., Srivastava, S., & Saggu, A. K. (2024, April). Revolutionising Arrhythmia Classification: Unleashing the Power of Machine Learning and Data Amplification for Precision Healthcare. In *2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT)* (pp. 516-522). IEEE.

[135] Kumar, R., Das, R., Garg, P., & Pandita, N. (2024, April). Duplicate Node Detection Method for Wireless Sensors. In *2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT)* (pp. 512-515). IEEE.

[136] Bhardwaj, H., Das, R., Garg, P., & Kumar, R. (2024, April). Handwritten Text Recognition Using Deep Learning. In *2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT)* (pp. 506-511). IEEE.

[137] Gill, A., Jain, D., Sharma, J., Kumar, A., & Garg, P. (2024, May). Deep learning approach for facial identification for online transactions. In *2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP)* (pp. 715-722). IEEE.

[138] Mittal, H. K., Dalal, P., Garg, P., & Joon, R. (2024, May). Forecasting Pollution Trends: Comparing Linear, Logistic Regression, and Neural Networks. In *2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP)* (pp. 411-419). IEEE.

[139] Malik, T., Nandal, V., & Garg, P. (2024, May). Deep Learning-Based Classification of Diabetic Retinopathy: Leveraging the Power of VGG-19. In *2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP)* (pp. 645-651). IEEE.

[140] Srivastava, A. K., Verma, I., & Garg, P. (2024, May). Improvements in Recommendation Systems Using Graph Neural Networks. In *2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP)* (pp. 668-672). IEEE.

[141] Aggarwal, A., Jain, D., Gupta, A., & Garg, P. (2024, May). Analysis and Prediction of Customer Churn and Retention Rates in the Telecom Industry Using Logistic Regression. In *2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP)* (pp. 723-727). IEEE.

[142] Mittal, H. K., Arsalan, M., & Garg, P. (2024, May). A Novel Deep Learning Model for Effective Story Point Estimation in Agile Software Development. In *2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP)* (pp. 404-410). IEEE.

[143] Shukla, S. M., Magoo, C., & Garg, P. (2024, November). Comparing Fine-Tuned LMs for Detecting LLM-Generated Text. In *2024, the 3rd Edition of IEEE Delhi Section Flagship Conference (DELCON)* (pp. 1-8). IEEE.

[144] Kumar, B., IQBAL, M., Parmer, R., Garg, P., Rani, S., & Agrawal, A. (2025, March). The Role of AI in Optimising Healthcare Appointment Scheduling. In *2025, the 3rd International Conference on Disruptive Technologies (ICDT)* (pp. 881-887). IEEE.

[145] Kumar, B., Garg, V., Ahmed, K., Garg, P., Choudhary, S., & Baniya, P. (2025, March). Enhancing Healthcare with Blockchain: Innovations in Data Privacy, Security, and Interoperability. In *2025, the 3rd International Conference on Disruptive Technologies (ICDT)* (pp. 932-938). IEEE.

[146] Raj, V., Prakash, B. K., Kumar, A., & Garg, P. (2024, December). Optimise the Time a Mercedes-Benz Spends on the Test Bench Using Stacking Ensemble Learning. In *2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS)* (pp. 445-450). IEEE.

[147] Kaushik, N., Kumar, H., Raj, V., & Garg, P. (2024, December). Proactive Fault Prediction in Microservices Applications Using Trace Logs and Monitoring Metrics. In *2024 International*

Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 410-415). IEEE.

[148] Kumar, A. A., Sri, C. V., Bohara, K. S. K., Setia, S., & Garg, P. (2024, December). Capnivesh: Financing Platform for Startups. In *2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS)* (pp. 261-265). IEEE.

[149] Bhandari, P., Setia, S., Kumar, K., & Garg, P. (2024, December). Optimising Cross-Platform Development with CI/CD and Containerization: A Review. In *2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS)* (pp. 175-180). IEEE.

[150] Chaudhary, A., & Garg, P. (2014). Detecting and diagnosing a disease using a patient monitoring system. *International Journal of Mechanical Engineering And Information Technology*, *2*(6), 493-499.