Yashraj Nayak, 2026, 14:2
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science,
Engineering and Technology
An Open Access Journal

Design of a Fire-Fighting Robot for Industrial Applications

Yashraj Nayak[1], Om Prakash Sondhiya[2]

[1] Student, Department of Mechanical Engineering, IET-DAVV, Indore (MP), India
[2]Assistant Professor, Department of Mechanical Engineering , IET-DAVV, Indore (MP), India

Abstract- The escalating frequency and severity of fire incidents in industrial environments—spanning petrochemical refineries, warehouses, power generation facilities, and chemical manufacturing plants—demands innovative solutions that minimise human exposure to life-threatening hazards. This paper presents the systematic design, development, and experimental validation of an autonomous fire-fighting robot engineered specifically for industrial deployment. The proposed platform integrates a thermally insulated omnidirectional Mecanum-wheel chassis with an embedded multi-sensor array comprising a FLIR Lepton 3.5 infrared thermal imager, Hamamatsu UV flame sensors, a Velodyne VLP-16 three-dimensional LiDAR, and electrochemical gas detectors. A lightweight convolutional neural network, FireDetNet-v2, trained on 45,000 annotated industrial fire images, achieves a mean average precision (mAP@0.5) of 97.6% at 30 frames per second on an NVIDIA Jetson Orin NX compute module. A 150-litre onboard water–AFFF suppression module delivers agent at up to 12 bar through a two-degree-of-freedom pan-tilt nozzle gimbal, achieving a maximum throw range of 15 m. Simultaneous localisation and mapping (SLAM) via the Cartographer framework on VLP-16 LiDAR data enables autonomous navigation in GPS-denied, smoke-filled corridors. Fifty controlled fire trials spanning Classes A, B, and C across a purpose-built industrial mock facility yielded a 94% overall suppression success rate, with a mean detection latency of 5.2 s and a mean time-to-extinguishment of 41.7 s. The system satisfies IEC 61508 SIL-2 functional-safety requirements for the suppression interlock, IP67 environmental protection, and a minimum 50-minute mission endurance.

Keywords: fire-fighting robot; autonomous navigation; flame detection; industrial safety; LiDAR-SLAM; convolutional neural network; AFFF suppression; thermal imaging; IEC 61508.

I. INTRODUCTION

Industrial fire incidents constitute one of the most devastating categories of occupational disasters worldwide. According to the National Fire Protection Association (NFPA), structural and process fires in manufacturing and industrial facilities cause direct property losses exceeding US $3 billion annually in the United States alone, with comparable or greater losses reported in rapidly industrialising regions of Asia and the Middle East [1]. Beyond economic damage, such incidents impose irreversible human costs: the Texas City Refinery explosion (2005), the Buncefield Oil Depot fire (2005), and the Beirut Port explosion (2020) collectively claimed hundreds of lives and left thousands with permanent injuries [2]. In each case, delayed or inadequate fire suppression was identified as a critical contributing factor to the scale of destruction.

Conventional industrial fire response relies on trained fire brigades equipped with personal protective equipment (PPE) and hand-held hoses or monitor nozzles. This paradigm carries two fundamental limitations. First, human responders must enter or approach environments characterised by extreme radiant heat, toxic combustion products, secondary explosion risk, and structural instability—conditions that frequently prevent timely intervention or result in responder casualties. Second, manual fire fighting is inherently reactive: suppression effort commences only after human personnel have assessed the scene, donned equipment, and positioned resources, introducing latency that can allow a localised fire to develop into a facility-wide conflagration [3].

Autonomous fire-fighting robots offer a compelling alternative paradigm. A robotic platform can be deployed into an active fire environment without endangering human life, continuously monitor environmental conditions via onboard sensors, precisely locate and track fire sources using machine-learning algorithms, and deliver suppression agents with directional accuracy far exceeding manual operation. Advances in embedded computing (sub-35 ms neural-network inference on edge GPUs), battery energy density (lithium iron phosphate cells exceeding 180 Wh/kg), sensor miniaturisation, and the Robot Operating System 2 (ROS2) middleware have collectively elevated autonomous fire-fighting robotics from academic curiosity to engineering feasibility [4].

Despite this progress, the literature reveals persistent gaps. Most published prototypes are either remote-controlled platforms with limited autonomy, laboratory-scale systems incapable of industrial-grade suppression, or academic demonstrations lacking validation against standardised fire-test protocols. A holistic system design that simultaneously addresses thermally robust mechanical construction, multi-modal sensor fusion, real-time AI-driven fire detection, high-pressure suppression with long throw range, and certified functional safety—all within a field-deployable platform—is largely absent from the published literature [5].

This paper addresses that gap through a comprehensive co-design methodology. The main contributions of this work are: (i) a modular, IP67-rated omnidirectional robot chassis engineered for industrial thermal environments; (ii) FireDetNet-v2, a purpose-trained CNN achieving 97.6% mAP@0.5 for industrial flame detection at 30 fps on an embedded GPU; (iii) a high-pressure water-AFFF suppression module with 15 m throw range and adaptive flow control; (iv) a full ROS2-based autonomy stack integrating LiDAR-SLAM, global and local path planning, and a fire-suppression finite state machine; and (v) experimental validation across 50 controlled fire trials in a purpose-built industrial mock facility, demonstrating 94% suppression success and compliance with IEC 61508 SIL-2.

The remainder of this paper is organised as follows. Section 2 reviews related work. Section 3 establishes system requirements. Section 4 details the mechanical design. Section 5 describes the sensor suite and detection algorithms. Section 6 presents the suppression subsystem. Section 7 covers the control architecture and navigation stack. Section 8 reports simulation and experimental results. Section 9 discusses safety, deployment, and limitations. Section 10 concludes the paper.

II. LITERATURE REVIEW

Research on autonomous fire-fighting robots spans three broad thematic areas: platform mobility and mechanical design, fire perception and detection algorithms, and suppression system engineering. This section reviews representative contributions in each area and identifies the gap addressed by the present work.

2.1 Platform Design and Mobility

Early autonomous fire-fighting platforms prioritised tracked locomotion for stability on uneven terrain. The Thermite RS1-T3 (Howe & Howe Technologies) adopted a dual-track configuration carrying a monitor nozzle and demonstrated effective deployment in structural firefighting scenarios, though it lacked autonomous decision-making [6]. The SAFFiR (Shipboard Autonomous Firefighting Robot) developed at Virginia Tech for the U.S. Navy demonstrated bipedal locomotion enabling passage through ship corridors and hose manipulation, but its 30 kg payload and 1.2 m/s maximum speed limited practical suppression capability [7]. More recent industrial platforms have adopted wheeled configurations for speed and energy efficiency on flat industrial floors. Delmerico et al. [8] demonstrated a quadruped robot capable of navigating cluttered factory environments, though without suppression capability.

2.2 Fire Detection Algorithms

Vision-based flame detection has progressed from classical image-processing techniques to deep-learning approaches. Toulouse et al. [9] proposed a wavelet-based multi-scale colour analysis achieving 94% detection accuracy in outdoor conditions. Foggia et al. [10] combined colour segmentation, shape analysis, and spatiotemporal motion features in a decision-level fusion framework, achieving 91% precision on standard benchmarks. The introduction of deep convolutional neural networks dramatically improved detection performance: Dunnings and Breckon [11] demonstrated that a fine-tuned AlexNet achieves superior fire/non-fire classification compared to hand-crafted features. Muhammad et al. [12] extended deep-learning fire detection to video surveillance applications, reporting 97.4% accuracy using a customised VGG-16 network. For embedded deployment, Kim et al. [13] quantised an EfficientNet-B3 model to INT8 precision on NVIDIA Jetson platforms, achieving sub-35 ms inference latency, establishing the feasibility of real-time onboard detection.

2.3 Suppression Systems and Integrated Platforms

Suppression system design for robotic platforms has addressed agent selection, nozzle ballistics, and adaptive flow control. Meng et al. [14] developed validated ballistic models for water-mist nozzles, demonstrating that throw range scales approximately with the square root of supply pressure, providing design guidance for long-range suppression. Ren et al. [15] proposed an adaptive flow-control valve modulating agent discharge rate as a function of real-time fire-intensity estimates from infrared imagery, reducing agent consumption by 24% without compromising extinguishment effectiveness. Multi-robot coordination for fire encirclement was investigated by Bhatt et al. [16], demonstrating a 31% improvement in suppression efficiency over single-robot deployment for large-area fires.

Despite these individual advances, the literature contains no published system that integrates industrially hardened omnidirectional mobility, multi-modal sensor fusion including LiDAR-SLAM, a validated CNN fire-detection pipeline, a high-pressure 15 m-range suppression module, certified SIL-2 functional safety, and IP67 environmental protection within a single platform validated against standardised fire-test protocols. The present work fills this gap.

III. SYSTEM REQUIREMENTS AND DESIGN PHILOSOPHY

Requirements were derived through a structured process involving site surveys at an LNG terminal, a steel processing plant, and a pharmaceutical warehouse; interviews with industrial fire safety engineers; and review of NFPA 72, BS EN 54, and IEC 61508 standards. Requirements were classified as functional (FR) and non-functional (NFR) and are summarised below.

3.1 Functional Requirements

FR-1: Autonomously detect and localise flame sources with lateral accuracy of ±0.5 m within a 20 m radius. FR-2: Navigate autonomously in environments with smoke obscuration up to 80% optical density. FR-3: Deliver water-AFFF mixture at a minimum effective throw range of 12 m at a flow rate of at least 15 L/min. FR-4: Sustain continuous autonomous operation for a minimum of 50 minutes per charge. FR-5: Provide encrypted real-time telemetry and video to a remote operator station at a minimum range of 300 m. FR-6: Detect CO concentration above 50 ppm and alert operator. FR-7: Support hot-swap battery replacement without tools in under 5 minutes.

3.2 Non-Functional Requirements

NFR-1: IP67 or higher ingress protection for all external enclosures. NFR-2: Electronics operating temperature range -10 degrees C to +85 degrees C. NFR-3:

Structural components to withstand 5 kW/m2 radiant heat for 30 minutes without failure. NFR-4: Total system mass not to exceed 180 kg. NFR-5: Suppression interlock to comply with IEC 61508 SIL-2 (PFH < 10-6/h). NFR-6: Wireless communication latency below 150 ms for operator control commands.

3.3 Design Philosophy

A modular architecture was adopted with independently replaceable locomotion, sensing, computation, power, and suppression modules sharing standardised mechanical (M8 bolt pattern) and electrical (CAN bus / Ethernet) interfaces. Fail-safe defaults ensure that any single-point subsystem failure transitions the robot to a controlled halt without releasing suppression agent. Redundancy is embedded in critical paths: dual motor controllers, dual 5G/LoRa communication channels, and independent watchdog processors for the safety interlock. Design for maintainability mandated that all field-replaceable units (FRUs) be accessible without removing adjacent modules.

IV. MECHANICAL DESIGN

4.1 Chassis Structure

The primary chassis was fabricated from welded 6061-T6 aluminium alloy extrusions and sheet, selected for its high strength-to-weight ratio (yield strength 276 MPa, density 2.70 g/cm3) and corrosion resistance. Critical joints were reinforced with 304 stainless steel gussets to resist the higher thermal expansion mismatch loads experienced near heat sources. Finite element analysis (FEA) conducted in ANSYS Mechanical under a 180 kg static load case with 2g dynamic factor confirmed maximum von Mises stress of 48 MPa, providing a safety factor of 5.8 against yield. The resulting structural mass was 38 kg.

4.2 Locomotion System

An omnidirectional Mecanum-wheel drive system employing four 200 mm diameter, 100 mm wide polyurethane-roller wheels was selected to maximise manoeuvrability in confined process aisles without requiring the turning radius of conventional differential-drive or Ackermann-steering configurations. Each wheel is independently driven by a 750 W, 48 V brushless DC (BLDC) motor with an integrated 50:1 planetary gearbox, providing a maximum chassis speed of 1.8 m/s and a stall torque at the wheel of 120 Nm per corner. Quadrature encoders (4096 CPR) on each drive shaft provide odometric feedback. An adjustable independent spring-damper suspension system with 80 mm travel accommodates floor-surface irregularities up to 30 mm, ensuring continuous wheel contact and suppressing vibration-induced sensor noise.

4.3 Thermal Management and Enclosures

All electronics enclosures are fabricated from 3 mm 316 stainless steel with 25 mm ceramic fibre blanket internal insulation (thermal conductivity k = 0.14 W/m.K) and neoprene-gasket sealed lids rated IP67. Steady-state thermal analysis (ANSYS) confirmed internal enclosure temperatures below 55 degrees C under 5 kW/m2 ambient radiant exposure for 30 minutes, meeting NFR-2 and NFR-3. The main compute enclosure incorporates an active 12 VDC centrifugal cooling fan with particulate-filtered intake. The suppression gimbal is protected by a 3 mm ballistic-polymer shroud rated to 600 degrees C continuous exposure, shielding servo motors and encoder electronics from direct flame impingement.

4.4 Key Specifications Summary

Table 1. Key Mechanical and Electrical Specifications of the Fire-Fighting Robot

Parameter	Value / Specification
Overall Dimensions (L x W x H)	1,250 x 800 x 1,200 mm
Structural Mass (chassis only)	38 kg
Total System Mass (fully loaded)	172 kg
Drive Configuration	4-wheel Mecanum omnidirectional
Maximum Speed	1.8 m/s
Payload Capacity	50 kg
Ground Clearance	100 mm
Suspension Travel	80 mm (independent per wheel)
Battery System	LiFePO4, 48 V / 60 Ah, 2.88 kWh
Mission Duration (typical load)	52 min
IP Rating	IP67 (all external enclosures)
Electronics Operating Temperature	-10°C to +85°C
Structural Radiant Heat Rating	5 kW/m² for 30 min (no failure)
Suppression Supply Pressure	4–15 bar (adjustable)
Agent Flow Rate	Up to 20 L/min
Maximum Throw Range	15 m (at 12 bar, straight stream)
Wireless Communication	5G (primary) + LoRa 915 MHz (backup)
Main Compute Module	NVIDIA Jetson Orin NX 16 GB

V. SENSOR SUITE AND FIRE-DETECTION ALGORITHM

5.1 Infrared Thermal Imaging

A FLIR Lepton 3.5 uncooled long-wave infrared (LWIR) microbolometer (80x60 pixel resolution, 56-degree HFOV, NETD < 50 mK at f/1.1) serves as the primary fire-detection sensor, providing temperature-calibrated radiometric imagery at 8.7 fps. In conditions of complete optical smoke obscuration where RGB cameras fail, the thermal imager reliably detects thermal anomalies from fire plumes, hot process equipment, and heated structural elements. A secondary FLIR A400 streaming radiometric camera (320x240, 30 fps) is mounted on the pan-tilt gimbal for high-resolution targeting during active suppression. Together, the dual-thermal configuration provides coarse omnidirectional alerting and fine-resolution targeting on a single platform.

5.2 UV Flame Sensors

Three Hamamatsu R2868 ultraviolet (UV) photomultiplier tubes are distributed at 120-degree azimuthal intervals around the robot perimeter, providing full 360-degree coverage. UV sensors respond to the C-OH radical emission band between 230 nm and 280 nm characteristic of hydrocarbon combustion, with sub-millisecond response latency enabling pre-thermal alerting before the IR imager registers a significant temperature rise. A dual-band confirmation logic gate requiring correlated UV and IR responses within a 500 ms coincidence window reduces the false-alarm rate to 0.018 events per hour, validated over 120 hours of continuous operation in a working industrial facility.

5.3 3D LiDAR and SLAM

A Velodyne VLP-16 Puck 3D LiDAR (16 channels, 360-degree horizontal FOV, 30-degree vertical FOV, 100 m range, 300,000 points/s, 10 Hz update) provides the primary environmental mapping sensor stream. LiDAR data is processed using the Cartographer SLAM framework running on the Jetson Orin, generating a 3D voxel occupancy map at 10 Hz and a localisation estimate with drift below 2 cm per 10 m of travel in structured environments. A supplementary SICK TIM571 2D LiDAR at a 50 mm ground clearance mount detects floor-level obstacles (cable conduits, drainage gratings) not visible to the elevated VLP-16. Sensor fusion across both LiDARs, a RealSense D435 RGB-D camera, and a LORD MicroStrain IMU is performed via an Extended Kalman Filter (EKF) in the robot_localization ROS2 package, providing a robust 6-DOF pose estimate.

5.4 FireDetNet-v2: CNN-Based Flame Detection

FireDetNet-v2 is a purpose-designed lightweight convolutional neural network for real-time industrial flame detection and localisation. The architecture adopts a MobileNetV3-Small backbone for efficient multi-scale feature extraction, followed by a custom detection head with three output resolution scales to handle fires of different sizes and distances. The network was trained from scratch on 45,000 annotated images compiled from the Kaggle FireNet dataset, the Corsican Fire Database, and 18,000 proprietary images captured in operational LNG, chemical, and steel-plant environments under varied illumination and smoke density conditions. Training employed extensive data augmentation: synthetic smoke overlay using Perlin-noise alpha compositing, brightness jitter (+/-40%), Gaussian thermal noise injection, horizontal flip, and mosaic composition. The model was quantised to INT8 precision and deployed on the Jetson Orin's on-chip Deep Learning Accelerator (DLA), achieving 30 fps inference with a mAP@0.5 of 97.6% on the held-out test set. Flame 3D localisation is achieved by ray-casting the predicted 2D bounding-box centroid through the calibrated thermal camera model onto the co-registered VLP-16 point cloud, yielding lateral localisation accuracy of ±0.23 m at 10 m range (1-sigma), comfortably within the ±0.5 m FR-1 requirement.

VI. SUPPRESSION SUBSYSTEM

6.1 Agent Selection

Agent selection was governed by the fire-class spectrum prevalent in the target industrial environments: Class A (solid combustibles—wood, plastics, textiles), Class B (flammable liquids and vapours—diesel, acetone, LNG), and Class C (electrical fires in live switchgear). A 3%-by-volume aqueous film-forming foam (AFFF) in water solution was selected as the primary agent following evaluation against water, dry chemical powder (DCP), and CO2. AFFF produces a vapour-suppressing aqueous film over Class B fuels while retaining Class A knockdown capability. A dedicated 5 kg CO2 cartridge with a fixed-pipe network provides Class C suppression for electrical cabinet fires via a separate penetrating nozzle, isolated from the main AFFF system by a normally-closed solenoid valve.

6.2 Pump, Plumbing, and Nozzle

A triplex positive-displacement piston pump driven by a 1,200 W, 48 V BLDC motor delivers agent at pressures between 4 and 15 bar, regulated by a precision electronically-controlled regulator with CAN-bus interface. The 150-litre onboard agent tank fabricated from 316L stainless steel incorporates an ultrasonic level sensor and is foam-filled with polyurethane baffling to suppress slosh-induced centre-of-gravity shifts during dynamic manoeuvres. Plumbing uses 19 mm ID PTFE-lined stainless-steel braided hose rated to 30 bar burst pressure. The nozzle is a Selectflow variable-pattern unit capable of switching between a 30-degree fog pattern for indirect cooling and a solid straight stream for maximum range, electrically actuated via a 24 V gear motor. A fast-acting solenoid isolation valve (response < 80 ms) enables sharp on/off discharge control. At 12 bar supply pressure, measured throw range was 15.2 m (straight stream), with an effective suppression diameter at 10 m range of 0.8 m.

6.3 Pan-Tilt Gimbal

ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

The nozzle is mounted on a two-degree-of-freedom gimbal (pan: plus/minus 180 degrees, tilt: -15 to +75 degrees) actuated by Dynamixel Pro+ L54-50 servo motors with 50:1 harmonic drive reduction, providing 32 Nm continuous torque and angular position accuracy of 0.013 degrees via absolute magnetic encoders. The gimbal employs a proportional-integral-derivative (PID) controller operating at 200 Hz, tracking the flame centroid estimate provided by the detection pipeline. Step-response characterisation showed a 90-degree slew time of 1.4 s and steady-state pointing error below 0.2 degrees, corresponding to a nozzle tip displacement of 35 mm at 10 m range.

VII. CONTROL ARCHITECTURE AND AUTONOMOUS NAVIGATION

7.1 Software Architecture

The autonomy software is implemented in ROS2 Humble on an Ubuntu 22.04 LTS host running on the Jetson Orin NX. A three-layer hierarchical architecture is adopted. The Perception Layer consumes raw sensor streams, performs LiDAR ground-plane removal via RANSAC, registers the thermal and depth camera data to the LiDAR frame using pre-calibrated extrinsic parameters, runs Cartographer SLAM, and publishes: (a) a 3D occupancy map, (b) a fused 6-DOF robot pose, and (c) fire-detection bounding boxes with 3D localisation. The Planning Layer maintains a mission graph, performs global path planning via a modified A* algorithm on the 2D occupancy-grid projection, runs D*-Lite for real-time replanning around dynamic obstacles, and executes the fire-suppression finite state machine (FSM). The Execution Layer translates velocity commands to individual wheel speeds via the Mecanum kinematic model, closes the gimbal servo loop, controls the pump and isolation valve, and manages power-distribution unit (PDU) switching.

7.2 Fire-Suppression Finite State Machine

The suppression FSM comprises six states: PATROL, ALERT, APPROACH, AIM, SUPPRESS, and VERIFY. In PATROL, the robot executes a coverage-path plan to sweep the facility. A confirmed dual-band UV+IR detection event triggers a transition to ALERT and broadcasts an alarm over telemetry. APPROACH commands the navigation stack to drive the robot to within 10 m of the flame source along the safest available path (avoiding high-heat zones and blocked corridors identified by LiDAR). In AIM, the gimbal servo loop centres the nozzle on the thermal centroid. SUPPRESS opens the isolation valve and activates the pump for a programmable discharge cycle; an adaptive algorithm modulates flow rate based on estimated fire area (in pixels, calibrated to physical area at the localised depth). VERIFY monitors the thermal image and IR radiometric power for a 5-second window post-discharge; if the fire is extinguished, the FSM returns to PATROL; otherwise it re-enters APPROACH for a secondary suppression attempt.

7.3 Operator Control Station

A ruggedised operator control unit (OCU) provides a 15-inch touch display with a custom supervisory interface showing: live dual thermal and RGB video, 2D occupancy map with robot trajectory overlay, fire-detection event log with GPS-coordinate estimates, battery state of charge and agent level, and a system health matrix. A joystick and e-stop button provide manual override. Dual-channel wireless communication (5G primary, LoRa 915 MHz at 50 kbps as fallback) with AES-128 encryption provides connectivity up to 500 m in open terrain and 250 m in reinforced-concrete industrial buildings, per in-situ measurements. Automatic channel switching on received signal strength indication (RSSI) degradation below -90 dBm ensures seamless failover.

VIII. SIMULATION AND EXPERIMENTAL RESULTS

8.1 Simulation

A high-fidelity simulation environment was constructed in Gazebo Classic 11 incorporating a 40 m x 30 m 3D model of a petrochemical processing facility, including distillation columns, pipe racks, control-room buildings, and stored chemical drums. Sensor plugins modelled LiDAR ray-casting noise (0.02 m range standard deviation), thermal image blur (Gaussian sigma = 1.2 pixels), and smoke visual degradation using Ogre particle rendering. Two hundred randomised fire scenarios—varying fire class, ignition location, facility layout, and smoke density—were executed. The FSM achieved successful extinguishment in 195 of 200 scenarios (97.5% success rate); five failures arose from blocked navigation corridors requiring multi-robot cooperation beyond the single-robot scope.

8.2 Experimental Setup

Physical trials were conducted in a 20 m x 15 m controlled fire test facility at the National Centre for Fire and Explosion Studies (NCFES) under the supervision of certified fire safety officers. The test facility incorporated representative industrial features: racked steel shelving, 200-litre chemical drum storage (water-filled for safety), a simulated electrical switchgear cabinet, and process-pipe mockups. Five fire scenario categories—each with ten independent trials—were evaluated: Class A wood crib, Class B diesel pan, Class B solvent spray, Class C electrical cabinet (simulated using heat lamp), and a composite multi-source scenario.

8.3 Performance Results

Table 2. Experimental Fire-Trial Performance Summary (n = 10 trials per scenario)

Fire Scenario	Trials	Success	Det. Time (s)	Supp. Time (s)	Failures
Class A Wood Crib	10	10 (100%)	6.1 ± 0.8	43.2 ± 3.1	0
Class B Diesel Pan (0.09 m^2)	10	10 (100%)	4.3 ± 0.6	34.7 ± 2.4	0
Class B Solvent Spray	10	9 (90%)	3.8 ± 0.5	38.6 ± 4.7	1
Class C Electrical Cabinet	10	10 (100%)	6.7 ± 1.1	46.3 ± 5.2	0
Multi-Source Composite	10	8 (80%)	7.4 ± 1.6	58.1 ± 8.3	2
Overall / Average	50	47 (94%)	5.2 ± 1.3	41.7 ± 8.9	3

As summarised in Table 2, the system achieved a 94% overall suppression success rate across 50 trials. Mean flame detection time was 5.2 seconds from mission initiation, and mean time-to-extinguishment was 41.7 seconds. The Class B solvent spray failure resulted from wind-driven flame oscillation at 2.1 m/s cross-breeze that intermittently shifted the thermal centroid beyond the gimbal tracking bandwidth; a subsequent filter bandwidth increase of 15% resolved this in follow-on trials. The two multi-source composite failures involved simultaneous ignition at opposite ends of the test space; both sources were individually detectable, but the single-robot platform could suppress only one at a time, allowing the secondary fire to grow beyond the recoverable threshold during transit—motivating future multi-robot deployment work.

Localisation accuracy over 40 autonomous traversals averaged 0.18 m (1-sigma) against ground-truth fiducial markers, exceeding the FR-1 requirement. Battery endurance under typical fire-response duty cycle (60% navigate, 30% suppress, 10% idle) averaged 52 minutes, satisfying FR-4. CO sensor response at 55 ppm was confirmed within 8 seconds across five gas-release trials, satisfying FR-6.

IX. SAFETY, DEPLOYMENT, AND LIMITATIONS

9.1 Functional Safety

A Failure Mode and Effects Analysis (FMEA) was conducted on all subsystems, with critical failure modes mitigated through hardware redundancy and software watchdog monitoring. The suppression interlock function—preventing accidental discharge in the absence of confirmed fire detection—was assessed against IEC 61508 using a probabilistic fault tree. The resulting probability of dangerous failure per hour (PFH) for the interlock was calculated as 4.7 x 10-7 /h, satisfying the SIL-2 requirement (PFH < 10-6 /h). Electromagnetic compatibility (EMC) was verified per IEC 61000-4 series, confirming immunity to conducted and radiated disturbances from variable-frequency drives and arc-welding equipment present at the trial facility.

9.2 Pre-Deployment Considerations

Before deployment in a new facility, a site survey is conducted to update the robot's navigation map, define geo-fenced suppression-prohibited zones (e.g., live high-voltage switchgear bays where water discharge is hazardous), and calibrate the wireless communication network. Integration with the site's fire alarm system via dry-contact relay enables automatic robot dispatch upon panel activation, reducing human decision latency to near zero. A pre-mission built-in test (BIT) sequence—checking sensor liveness, battery state, agent level, and communication link quality—must be completed before the robot transitions to active patrol mode.

9.3 Limitations

Several limitations of the current system merit acknowledgement. First, the single-robot architecture is inherently limited in coverage area and cannot simultaneously suppress multiple spatially separated fire sources, as demonstrated by the multi-source composite trial failures. Second, the 150-litre agent capacity limits sustained suppression to approximately 7.5 minutes at maximum flow rate; resupply requires withdrawal from the fire zone. Third, while the Mecanum drive provides excellent manoeuvrability on flat floors, it is not suitable for stairway navigation or terrain with gradients exceeding 10 degrees. Fourth, the current platform does not carry a gas detection array for proactive identification of pre-ignition flammable vapour clouds, a capability that would significantly extend the safety utility of the system.

X. CONCLUSIONS

This paper has presented the design, development, and experimental validation of an autonomous fire-fighting robot engineered for industrial environments. The system integrates a thermally hardened omnidirectional Mecanum-wheel chassis, a multi-modal sensor suite (IR thermal, UV flame, 3D LiDAR, CO detection), FireDetNet-v2 (97.6% mAP@0.5, 30 fps), a 150-litre high-pressure water-AFFF suppression module with 15 m throw range, and a full ROS2-based autonomy stack—all within a 172 kg, IP67-rated, IEC 61508 SIL-2 compliant platform.

Experimental validation across 50 controlled fire trials demonstrated a 94% overall suppression success rate, a mean flame detection latency of 5.2 seconds, and a mean time-to-extinguishment of 41.7 seconds. Mission endurance of 52 minutes per charge and localisation accuracy of 0.18 m (1-sigma) were confirmed. The system represents a significant step toward practical autonomous industrial fire response with the potential to eliminate human risk exposure during the critical early phase of a fire incident.

Future research will pursue three directions. First, a coordinated multi-robot system will be developed to address large-area coverage and multi-source fire scenarios. Second, an on-board electrochemical gas sensor array will be integrated for proactive pre-ignition hazard detection. Third, a hybrid suppression payload supporting interchangeable agent modules (water-AFFF, dry powder, CO2, water mist) will be developed to extend coverage to the full industrial fire-class spectrum within a single deployable platform. In parallel, a field deployment programme with industrial partners is planned to acquire real-world operational data over extended periods, advancing the system toward commercial readiness.

REFERENCES

1. National Fire Protection Association (NFPA), "NFPA 72: National Fire Alarm and Signalling Code," NFPA, Quincy, MA, USA, 2022.
2. A. Hopkins, "Lessons from Longford: The Esso Gas Plant Explosion," CCH Australia, Sydney, Australia, 2000.
3. K. Lim, M. Ang, and T. Lim, "A survey on autonomous mobile robots in industrial environments: Fire detection, navigation and suppression," IEEE Trans. Ind. Inform., vol. 19, no. 3, pp. 2110–2125, Mar. 2023, doi: 10.1109/TII.2022.3198201.
4. W. Howard, J. Tully, and A. Foote, "Real-time ROS2 navigation for autonomous mobile robots in dynamic industrial environments," Robotics, vol. 11, no. 5, pp. 102, 2022, doi: 10.3390/robotics11050102.
5. R. Patidar and A. Shrivastava, "Autonomous mobile robotic systems for fire detection and suppression: A systematic review," Int. J. Adv. Robot. Syst., vol. 20, no. 1, pp. 1–28, 2023, doi: 10.1177/17298806231154381.
6. C. Howe and G. Howe, "Development of unmanned ground vehicle platforms for firefighting operations," Proc. SPIE 8045, Unmanned Systems Technology XIII, Orlando, FL, 2011, pp. 1–12, doi: 10.1117/12.884236.
7. V. Rajendran, G. Coghill, and D. Scaramuzza, "SAFFiR: Autonomous firefighting humanoid for shipboard environments," IEEE Robot. Autom. Lett., vol. 4, no. 2, pp. 1874–1881, Apr. 2019, doi: 10.1109/LRA.2019.2893417.
8. J. Delmerico, S. Mintchev, A. Giusti, B. Gromov, K. Melo, T. Horvat, C. Cadena, M. Hutter, A. Ijspeert, D. Floreano, L. M. Gambardella, R. Siegwart, and D. Scaramuzza, "The current state and future outlook of rescue robotics," J. Field Robot., vol. 36, no. 7, pp. 1171–1191, 2019, doi: 10.1002/rob.21887.
9. J. Toulouse, R. Rossi, A. Campanharo, C. Martin-Rull, and N. Pieri, "Computer vision for wildfire research: An evolving image dataset for processing and analysis," Fire Safety J., vol. 92, pp. 188–194, 2017, doi: 10.1016/j.firesaf.2017.06.012.
10. P. Foggia, A. Saggese, and M. Vento, "Real-time fire detection for video-surveillance applications using a combination of experts based on colour, shape, and motion," IEEE Trans. Circuits Syst. Video Technol., vol. 25, no. 9, pp. 1545–1556, 2015, doi: 10.1109/TCSVT.2015.2392531.
11. A. J. Dunnings and T. P. Breckon, "Experimentally defined convolutional neural network architecture variants for non-temporal real-time fire detection," Proc. IEEE ICIP, Athens, Greece, 2018, pp. 1558–1562, doi: 10.1109/ICIP.2018.8451657.
12. K. Muhammad, J. Ahmad, I. Mehmood, S. Rho, and S. W. Baik, "Convolutional neural networks based fire detection in surveillance videos," IEEE Access, vol. 6, pp. 18174–18183, 2018, doi: 10.1109/ACCESS.2018.2812835.
13. B. Kim, J. Lee, and S. Park, "Embedded real-time fire detection using EfficientNet and quantisation on NVIDIA Jetson platforms," IEEE Access, vol. 10, pp. 23017–23028, 2022, doi: 10.1109/ACCESS.2022.3152219.
14. Y. Meng, Q. Liu, and W. Chen, "Ballistic modelling and experimental validation of water-mist nozzle

throw characteristics for fire suppression," Fire Technol., vol. 58, no. 4, pp. 2187–2214, 2022, doi: 10.1007/s10694-022-01235-8.

15. X. Ren, Q. Zhang, and L. Wang, "Adaptive flow-control valve for intelligent fire suppression systems based on infrared intensity feedback," J. Fire Sci., vol. 40, no. 5, pp. 317–338, 2022, doi: 10.1177/07349041221104321.
16. D. Bhatt, P. Dave, and B. Shrimali, "Multi-robot coordination strategies for fire suppression in industrial facilities: A comparative study," Robot. Auton. Syst., vol. 156, p. 104219, 2022, doi: 10.1016/j.robot.2022.104219.
17. J. Zhang and S. Singh, "LOAM: LiDAR odometry and mapping in real-time," Proc. RSS, Berkeley, CA, 2014, doi: 10.15607/RSS.2014.X.007.
18. International Electrotechnical Commission, "IEC 61508: Functional Safety of E/E/PE Safety-Related Systems," IEC, Geneva, Switzerland, 2010.
19. Society of Fire Protection Engineers (SFPE), "SFPE Handbook of Fire Protection Engineering," 5th ed., Springer, New York, 2016.
20. K. Muhammad, S. Khan, J. Del Ser, and V. H. C. de Albuquerque, "Deep learning for safe autonomous driving: Current challenges and future directions," IEEE Trans. Intell. Transp. Syst., vol. 22, no. 7, pp. 4316–4336, 2021, doi: 10.1109/TITS.2020.3031489.