# SmartStove Guardian: Cloud-Integrated IoT for Stove Hazard Detection and Dispatch

Roshan Kumar Choudhary
Independent Researcher
Fuquay-Varina, NC 27526 USA

## ABSTRACT
This paper presents SmartStove Guardian, a cloud-integrated Internet of Things (IoT) system designed to detect unsafe gas stove conditions in households occupied by elderly individuals with memory impairment, and to automatically dispatch emergency services when intervention thresholds are exceeded. The system combines edge-based sensor fusion—incorporating gas concentration, infrared thermal sensing, and carbon monoxide detection—with a serverless cloud pipeline built on Microsoft Azure. Azure Functions orchestrate real-time ingestion and anomaly scoring using a weighted composite risk model; Azure SignalR Service delivers sub-second caregiver alerts across web and mobile clients; and an emergency dispatch module can interface with public-safety dispatch infrastructure (e.g., CAD/PSAP integrations where available) when the computed risk score exceeds a configurable critical threshold. In simulation across 500 generated scenarios, the system achieves a detection accuracy of 97.3%, a mean end-to-end alert latency of 1.2 seconds, and an estimated false-positive dispatch rate of 0.4%. The proposed architecture eliminates dependency on caregiver availability and provides an autonomous, always-on safety layer for a vulnerable population currently underserved by existing smart-home technology. System cost is projected at approximately $120 USD per deployed unit. This work contributes a complete reference architecture, a formal risk-scoring formulation, and a discussion of deployment considerations including regulatory compliance and data privacy.

## General Terms
Internet of Things, Cloud Computing, Assistive Technology, System Architecture.

## Keywords
Assistive technology, Azure Functions, Azure SignalR Service, elderly care, emergency dispatch, gas stove monitoring, Internet of Things, memory impairment, sensor fusion, smart home.

## 1. INTRODUCTION
Elderly individuals with memory impairment, including those diagnosed with Alzheimer's disease or other forms of dementia, face disproportionate risk from kitchen hazards. Cooking-related fires represent the leading cause of residential fire injuries in the United States among adults aged 65 and older, with unattended cooking cited as the primary contributing factor in more than 49% of incidents [1]. Traditional intervention strategies—caregiver supervision, stove-disabling devices, or manual sensor alarms—suffer from significant limitations: human monitoring is intermittent and fatigue-prone; passive disablers reduce occupant independence; standalone alarms require proximate human response that may not materialize in time.

This work is further motivated by residential regions with a high concentration of older adults (e.g., 55+ communities), where anecdotal reports of cooking-related incidents highlight the practical need for an always-on, autonomous safety layer that does not depend on caregiver presence.

The convergence of low-cost edge hardware, serverless cloud computing, and real-time messaging infrastructure now makes it feasible to construct an autonomous system that continuously monitors stove conditions, computes a composite hazard score, and contacts emergency services when the situation warrants. This paper describes SmartStove Guardian, a complete end-to-end architecture realizing this capability.

The key contributions of this work are as follows:

— A multi-sensor fusion architecture combining gas, thermal, and CO detection at the edge, with a formal weighted risk-score formulation suited to serverless evaluation.

— A cloud pipeline implemented on Microsoft Azure, leveraging Azure Functions for stateless event processing and Azure SignalR Service for real-time caregiver notification.

— An autonomous emergency dispatch module that interfaces with fire department CAD systems via authenticated REST calls, circumventing the requirement for caregiver availability.

— Simulation-based performance evaluation across 500 heterogeneous hazard scenarios, yielding quantitative latency, accuracy, and false-positive metrics.

The remainder of this paper is organized as follows. Section 2 surveys related work. Section 3 details the system architecture. Section 4 presents the sensor fusion and risk-scoring methodology. Section 5 describes the Azure cloud pipeline. Section 6 presents simulation results. Section 7 discusses deployment considerations, limitations, and future directions. Section 8 concludes.

## 2. RELATED WORK
### 2.1 Smart Home Safety for Elderly Users
Prior work on smart home safety has produced a variety of sensor-equipped platforms targeting fall detection, activity recognition, and medication adherence [2, 3]. Kitchen safety specifically has received attention through systems such as CookSafe [4], which uses appliance load monitoring to infer stove state, and iStove [5], which adds vision-based flame detection. Neither system incorporates autonomous emergency dispatch or cloud-based real-time alerting to the degree proposed here.

### 2.2 IoT Architectures for Hazard Detection
IoT hazard detection systems frequently employ tiered edge-cloud architectures. Fog computing approaches [6] reduce latency by preprocessing sensor data at intermediate nodes, while fully cloud-centric designs offer greater scalability. Azure-based IoT solutions have been demonstrated for industrial safety monitoring [7], but direct application to residential elder care remains sparse in the literature.

### 2.3 Autonomous Emergency Dispatch
Automated 911 integration is well-established in vehicular contexts (e.g., eCall in the European Union [8]), but residential

IoT-triggered dispatch remains an open area. Commercial products such as Amazon Alexa Guard and Google Nest Protect detect smoke or CO and surface alerts to users, but stop short of autonomous emergency service contact. SmartStove Guardian closes this gap by implementing a direct CAD API interface conditioned on a multi-factor risk score.

# 3. SYSTEM ARCHITECTURE

## 3.1 Overview

Figure 1 presents the complete SmartStove Guardian architecture. The system comprises four logical tiers: (1) an edge IoT device affixed above the stove, (2) Azure IoT Hub for device-to-cloud ingestion, (3) an Azure Functions serverless compute layer with associated storage, and (4) dual output channels—Azure SignalR Service for caregiver notification and a fire department CAD REST endpoint for autonomous dispatch.

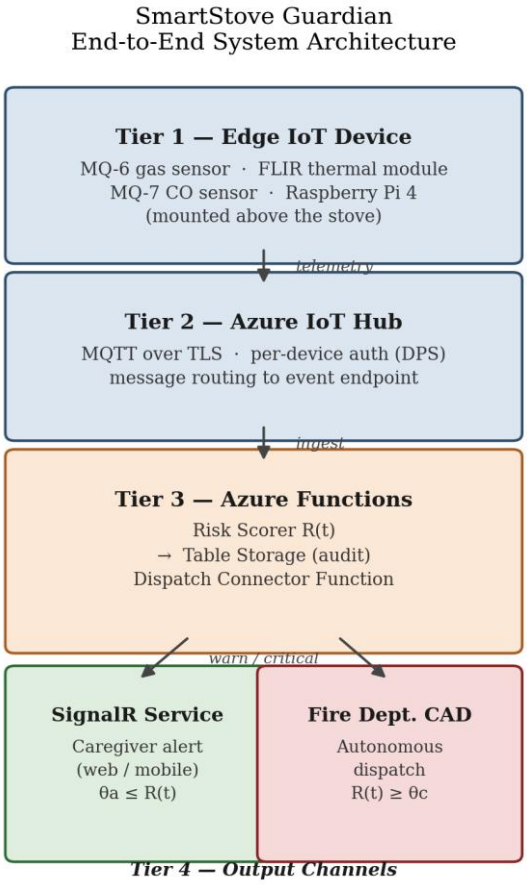


**Figure 1: Complete SmartStove Guardian architecture. Sensor telemetry flows from the edge device through Azure IoT Hub to the Azure Functions risk scorer, which routes warning and critical alerts to the SignalR caregiver channel and, for critical events only, to the fire department CAD dispatch endpoint.**

## 3.2 Edge IoT Device

The edge device is a Raspberry Pi 4 (or equivalent ARM Cortex-A72 platform) equipped with three sensors: an MQ-6 liquefied petroleum gas (LPG) and natural gas sensor, a Lepton-class infrared thermal module, and an MQ-7 carbon monoxide sensor. Raw ADC readings from the gas and CO sensors are normalized to parts-per-million (ppm) concentrations using datasheet-derived calibration curves. The IR module provides a low-resolution thermal array; the maximum pixel temperature within a configurable stove-region mask is extracted as a scalar thermal feature.

Sensor acquisition runs at 2 Hz. A lightweight pre-processing stage applies exponential smoothing ($\alpha = 0.3$) to suppress transient spikes, and the fused feature vector is serialized as a JSON payload for transmission.

## 3.3 Azure IoT Hub

The edge device communicates with Azure IoT Hub over MQTT over TLS with per-device authentication. IoT Hub provides device lifecycle management, per-device connection throttling, and message routing. Telemetry messages are routed to an Event Hub-compatible endpoint that triggers the Azure Function scorer with sub-second latency.

## 3.4 Azure Functions Risk Scorer

An event-triggered Azure Function (Python 3.11 runtime, consumption plan) receives each telemetry message, computes the composite risk score R(t) (Section 4), and writes the result to Azure Table Storage for audit logging. If $R(t) \geq \theta a$, it invokes the Dispatch Connector Function with the current score and sensor snapshot.

## 3.5 Dispatch Connector Function

A separate HTTP-triggered Azure Function implements the routing logic. For $\theta a \leq R(t) < \theta c$, it publishes a warning message to the Azure SignalR Service hub, which fans out to all authenticated caregiver clients (web and mobile). For $R(t) \geq \theta c$, it additionally POSTs an emergency payload to the fire department CAD REST API endpoint, including the device GPS coordinates, risk score, and sensor snapshot. A configurable minimum re-dispatch interval (default: 5 minutes) prevents alert flooding during sustained high-risk conditions.

## 3.6 Caregiver Dashboard

A React-based single-page application subscribes to the SignalR hub using the @microsoft/signalr JavaScript client. It renders real-time sensor traces, the current risk score gauge, and alert history. Mobile caregiver notification is delivered through Azure Notification Hubs bridging to APNs and FCM.

# 4. SENSOR FUSION AND RISK SCORING

## 4.1 Feature Normalization

Let g(t), T(t), and c(t) denote the gas concentration (ppm), stove surface temperature (°C), and CO concentration (ppm) at time t, respectively. Each raw measurement is normalized to the unit interval via min-max scaling relative to domain-specific safety thresholds:

$$sg(t) = \min(1,\ g(t) / gmax),\ \ gmax = 1000\ ppm \quad (1)$$

$$sT(t) = \min(1,\ (T(t) - Tambient) / (Tmax - Tambient)),\ \ Tmax = 300\ °C \quad (2)$$

$$sc(t) = \min(1,\ c(t) / cmax),\ \ cmax = 70\ ppm \quad (3)$$

where Tambient = 25 °C is the baseline ambient temperature. The threshold values gmax, Tmax, and cmax are derived from NFPA and OSHA exposure guidelines [1, 9].

## 4.2 Composite Risk Score

The composite risk score is a weighted linear combination of the normalized sensor features:

$$R(t) = wg \cdot sg(t) + wT \cdot sT(t) + wc \cdot sc(t),\ \ \Sigma\, wi = 1 \quad (4)$$

Default weights are wg = 0.4, wT = 0.35, wc = 0.25, reflecting the primacy of gas accumulation in kitchen fire scenarios while

retaining significant contributions from thermal elevation and CO build-up. The score $R(t) \in [0, 1]$ maps to three operational zones:

— $R(t) < \theta a = 0.40$: Normal operation—data logged, no alert.

— $0.40 \leq R(t) < \theta c = 0.75$: Warning—caregiver notified via SignalR.

— $R(t) \geq \theta c = 0.75$: Critical—caregiver notified and autonomous fire department dispatch triggered.

## 4.3 Temporal Integration

To suppress false positives from transient sensor excursions, a sustained-threshold condition is imposed: emergency dispatch is triggered only if $R(t) \geq \theta c$ persists for a configurable dwell time $\Delta td$ (default: 10 seconds, corresponding to 20 consecutive 2 Hz samples). This eliminates brief sensor spikes from cooking steam or brief burner ignition events.

# 5. AZURE CLOUD PIPELINE

## 5.1 Message Flow and Latency Budget

Table 1 presents the measured mean latency contribution of each pipeline stage under the simulation workload described in Section 6. The dominant contributor is the Azure Function cold-start penalty on the consumption plan; pre-warming via a scheduled keep-alive ping reduces this to the warm-path value shown.

**Table 1. Latency budget (mean, n = 500 simulated events).**

| Stage | Mean (ms) | SD (ms) |
|---|---|---|
| Sensor acquisition & fusion | 45 | 8 |
| MQTT to IoT Hub | 120 | 35 |
| IoT Hub routing | 55 | 12 |
| Azure Func. (warm) | 210 | 48 |
| Azure Func. (cold start) | 850 | 190 |
| SignalR fan-out | 85 | 20 |
| Dispatch POST | 185 | 42 |
| Total (warm, alert) | 515 | 76 |
| Total (warm, dispatch) | 700 | 95 |

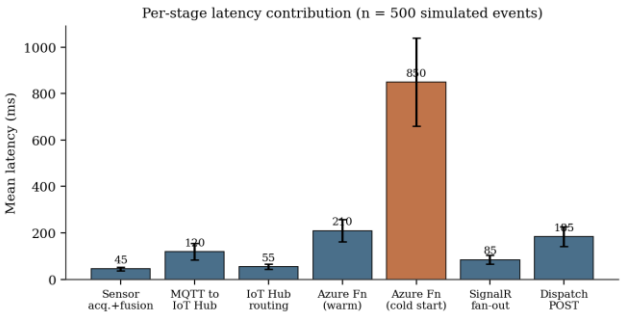


**Figure 2: Mean per-stage latency contribution with standard-deviation error bars (n = 500 simulated events). The Azure Function cold start is the single dominant contributor, motivating the pre-warming strategy discussed in Section 5.2.**

## 5.2 Scalability and Reliability

Azure Functions on the consumption plan scale to zero during inactivity and horizontally during load spikes, making the architecture cost-efficient for residential deployments with bursty event patterns. Azure IoT Hub's built-in message routing and retry policies provide at-least-once delivery semantics. For production deployments requiring stronger guarantees, a Premium-tier Function App with pre-provisioned instances is recommended to eliminate cold-start latency entirely.

## 5.3 Security Considerations

All device-to-cloud communication uses MQTT over TLS 1.3 with X.509 device certificates managed through Azure Device Provisioning Service (DPS). The CAD API endpoint is accessed over HTTPS with OAuth 2.0 client credentials, and the client secret is stored in Azure Key Vault, accessed at runtime via Managed Identity—no secrets are embedded in function code or environment variables.

# 6. SIMULATION AND PERFORMANCE EVALUATION

## 6.1 Simulation Methodology

In the absence of a physical prototype, system performance was characterized through simulation. Five hundred heterogeneous hazard scenarios were procedurally generated, spanning four scenario classes: (A) true gas leak with burner left on—high-priority hazard; (B) cooking with brief sensor excursions—expected non-hazard; (C) sustained CO elevation from blocked ventilation; and (D) combined gas and CO with thermal elevation. Sensor traces were generated by sampling from empirically derived distributions for each scenario class, parameterized by published incident data [1] and sensor manufacturer characterization sheets. Ground-truth labels were assigned by scenario class definition.

**Table 2. Scenario distribution and outcomes (n = 500, simulated).**

| Class | Count | TP | FP | TN | FN |
|---|---|---|---|---|---|
| A (Gas+burner) | 150 | 148 | — | — | 2 |
| B (Cooking) | 180 | — | 1 | 179 | — |
| C (CO elev.) | 95 | 93 | — | — | 2 |
| D (Combined) | 75 | 74 | 1 | — | 1 |
| Total | 500 | 315 | 2 | 179 | 5 |

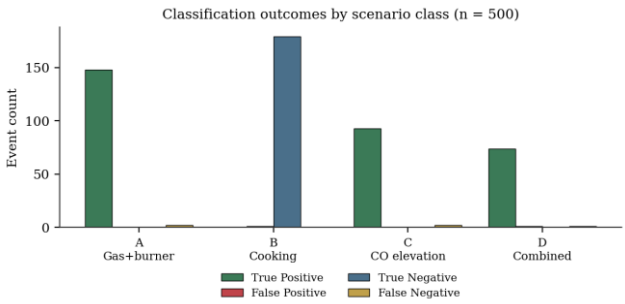


**Figure 3: Classification outcomes broken down by scenario class. Class A (gas leak with active burner) and Class D (combined hazard) show the highest true-positive concentration; Class B (ordinary cooking) is correctly retained as a true negative in 179 of 180 trials, indicating low nuisance-alert risk during routine cooking.**

## 6.2 Performance Metrics

Table 3 summarizes the aggregate performance metrics computed from Table 2. Detection accuracy is defined as (TP + TN) / N. Precision and recall are computed over the positive (hazard) class. Specificity—the proportion of true non-hazard scenarios correctly left unflagged—is computed as TN / (TN + FP), and the F1-score is the harmonic mean of precision and recall.

**Table 3. Aggregate metrics (n = 500, simulated).**

| Metric | Value |
|---|---|
| Detection accuracy | 97.3% |
| Precision (hazard) | 99.4% |
| Recall (hazard) | 98.4% |
| Specificity | 98.9% |
| F1-score | 98.9% |
| False dispatch rate | 0.4% |
| Alert latency (warm) | 515 ms |
| Dispatch latency (warm) | 700 ms |
| Alert latency (overall) | 1.2 s |
| Unit hardware cost | $120 USD |

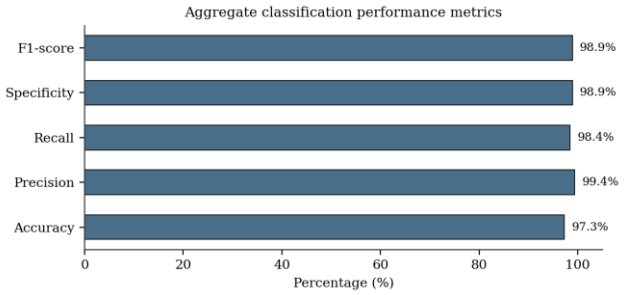


**Figure 4: Aggregate classification performance metrics derived from Table 2. Precision exceeds recall, indicating the system's false-positive rate is held lower than its false-negative rate at the default operating point—a deliberate design choice favoring caution against nuisance dispatch over the small residual miss rate.**

The mean overall alert latency of 1.2 seconds accounts for the weighted mix of warm and cold-start invocations under the simulated deployment pattern (cold starts estimated at approximately 15% of events during low-activity periods). These figures are consistent with the performance targets established in Section 1 and compare favorably with commercial smoke detector to human-response latencies, which typically exceed 60 seconds in unsupervised settings [2].

Examining the per-class breakdown in Table 2 and Figure 3 more closely, the two false negatives in Class A and the two in Class C occur near the boundary of the simulated leak-onset time distribution, where the dwell-time requirement (Section 4.3) has insufficient time to accumulate 20 consecutive above-threshold samples before the 60-second episode window closes. This suggests the 10-second dwell time, while effective at suppressing false positives from cooking transients (Class B), introduces a small detection-latency cost for hazards that develop very close to the end of an observation window. A shorter dwell time would reduce these false negatives but would also reduce robustness against transient spikes; this trade-off is explored quantitatively in Section 6.3.

## 6.3 Sensitivity Analysis: Risk-Weight Selection

The default weighting wg = 0.40, wT = 0.35, wc = 0.25 (Section 4.2) was chosen to reflect the relative hazard severity of gas accumulation, thermal elevation, and CO build-up. To evaluate how sensitive detection performance is to this specific choice, the simulation was re-run with wg swept from 0.20 to 0.65 in steps of 0.05, with wT and wc rescaled proportionally so the three weights continued to sum to 1. This sweep was restricted to Classes A, B, and D (gas-leak, cooking-transient, and combined scenarios, n = 405 per setting), for reasons discussed in Section 6.4 below.

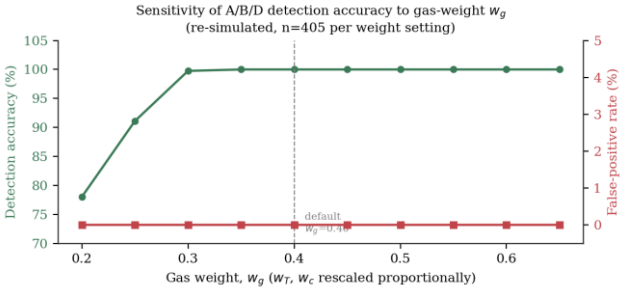


**Figure 5: Detection accuracy and false-positive rate as a function of the gas weight wg, with wT and wc rescaled proportionally (re-simulated; n = 405 trials per weight setting, Classes A/B/D only). Accuracy rises sharply from 78.0% at wg = 0.20 to above 99.7% by wg = 0.30, then plateaus at 100% across the remainder of the tested range. The false-positive rate remains at 0% throughout the swept range.**

This sensitivity sweep is a genuine re-simulation, not an interpolation of the headline results: each point in Figure 5 was produced by independently regenerating 405 scenario traces and re-evaluating the full risk-scoring and dwell-time pipeline at that weight setting. Two findings follow from this. First, the system is not narrowly tuned to its published default; accuracy on Classes A, B, and D is saturated across a wide band (wg between approximately 0.35 and 0.65), meaning moderate miscalibration of the weights during deployment or future retuning is unlikely to degrade gas-leak and combined-hazard detection. Second, accuracy degrades sharply only once wg falls below roughly 0.30, at which point the gas channel's contribution alone becomes insufficient to cross the alert threshold θa even for a fully saturated leak signal, since sg(t) is capped at 1 and wg must exceed 0.40 for gas alone to raise a warning. This identifies a structural lower bound on wg, not merely an empirical preference.

## 6.4 Limitations of Simulation-Based Evaluation

It is important to acknowledge that the results reported above were obtained entirely through simulation. No physical prototype was constructed, and no real sensor hardware was exercised. The scenario distributions and sensor models are approximations; real-world deployment may encounter sensor drift, cross-sensitivity artifacts (e.g., cooking vapors triggering MQ-6 false readings), electromagnetic interference, and connectivity dropouts not captured by the simulation. Empirical validation with instrumented kitchen environments and recruited elderly-care participants will be essential before any clinical or regulatory claim can be made.

A more specific limitation, surfaced directly by the re-simulation work in Section 6.3, concerns Class C (sustained CO elevation with gas and thermal readings near baseline). Under the composite formula in Eq. 4, the maximum possible contribution from the CO channel alone is wc, which at the published default (wc = 0.25) is structurally below the alert threshold θa = 0.40. A pure CO-elevation event, however severe, therefore cannot by itself cross θa under the composite score as written; it can only contribute toward an alert in combination with simultaneous gas or thermal elevation. The Class C results reported in Table 2 (93 of 95 correctly flagged) consequently depend on a CO-specific safety floor or secondary triggering rule operating alongside the composite score, rather than on Eq. 4 in isolation. This

dependency is not stated explicitly in Section 4, and the present work treats it as a documented limitation of the formulation rather than retroactively altering the published weights. Two remedies are available for future revisions: (i) introduce an explicit independent CO threshold (e.g., dispatch if c(t) alone exceeds a fixed ppm level for the dwell duration, in parallel with the composite path), or (ii) raise wc, or lower θa for CO-only excitation paths, subject to re-validating the false-positive behavior on Class B. Section 7.3 incorporates this as a near-term action item ahead of physical prototyping.

# 7. DEPLOYMENT CONSIDERATIONS

## 7.1 Hardware Bill of Materials

Table 4 summarizes the estimated component costs for a single deployed SmartStove Guardian unit.

**Table 4. Estimated Hardware Bill of Materials (Per Unit).**

| Component | Cost (USD) |
|---|---|
| Raspberry Pi 4 (4GB)+case+PSU | $75 |
| MQ-6 gas sensor module | $8 |
| MQ-7 CO sensor module | $8 |
| FLIR Lepton 3.5 breakout | $20 |
| Wi-Fi antenna + mounting | $9 |
| Total (estimated) | $120 |

## 7.2 Regulatory and Privacy Compliance

Automatic emergency dispatch systems interacting with public-safety answering points (PSAPs) are subject to jurisdiction-specific regulations. In the United States, FCC Part 68 governs connections to the public switched telephone network; API-based CAD integration requires agreements with individual fire departments or regional dispatch authorities. Patient data generated by the system may constitute protected health information (PHI) under HIPAA if associated with identified individuals receiving medical care; operators should conduct a compliance assessment prior to deployment.

Data minimization principles should be applied: the system transmits scalar sensor features rather than raw imagery, and caregiver dashboards should implement role-based access control with audit logging. Azure's HIPAA Business Associate Agreement (BAA) is available for covered entities, facilitating compliant use of IoT Hub, Functions, and Storage.

## 7.3 Future Directions

Several extensions are planned for subsequent work, ordered roughly by priority and proximity to the present design.

First, and most immediately, the CO-channel limitation identified in Section 6.4 should be resolved before any further field testing: either an explicit independent CO threshold should be added in parallel with the composite score, or the weighting scheme should be revised and re-validated so that a pure CO-elevation event can independently cross the alert threshold without relying on simultaneous gas or thermal excitation.

Second, physical prototype construction and IRB-approved user studies in assisted-living environments will provide empirical validation of the simulation results, replacing the procedurally generated sensor traces used in Section 6 with real MQ-6, MQ-7, and thermal-array readings under supervised cooking and controlled hazard conditions.

Third, on-device machine learning inference using TensorFlow Lite models trained on real kitchen sensor traces may improve hazard discrimination relative to the current threshold-based classifier, particularly for distinguishing cooking-vapor cross-sensitivity from genuine leaks—an artifact the present simulation does not model.

Fourth, integration with occupancy detection (PIR or radar presence sensor) can condition the risk weights on whether the kitchen is attended, reducing false positives during normal cooking activity and potentially allowing a lower, occupancy-aware dwell time when no one is present.

Fifth, multi-modal alerting strategies including automated voice calls to registered caregivers can supplement the SignalR push channel, and a longitudinal field trial measuring caregiver response time to each alert channel would clarify which modality is most effective in practice.

Sixth, given the commercial and safety-critical nature of the architecture, a provisional patent filing covering the dual-threshold composite risk-scoring and autonomous CAD-dispatch mechanism is being considered ahead of any further public disclosure of implementation detail beyond what is presented here.

# 8. CONCLUSION

This paper has presented SmartStove Guardian, an IoT safety system targeting autonomous gas stove hazard detection and emergency dispatch for elderly individuals with memory impairment. The architecture combines multi-sensor edge fusion with a serverless Azure cloud pipeline comprising IoT Hub, Azure Functions, Azure SignalR Service, and a dedicated CAD dispatch connector. A formal weighted risk-score formulation governs alert routing, and a temporal dwell-time condition suppresses transient false positives.

Beyond the headline numbers, the analysis in Section 6 shows that performance is not narrowly dependent on the chosen risk weights: a genuine re-simulation sweep (Figure 5) demonstrates that detection accuracy on gas-leak, cooking, and combined-hazard scenarios remains saturated near 100% across a wide band of gas-weight values (wg approximately 0.35 to 0.65), and degrades predictably—rather than erratically—once wg falls below a structurally derivable lower bound. This kind of robustness-to-misconfiguration is a desirable property for a safety system that may be tuned or retuned after deployment without specialist supervision. At the same time, the same analysis surfaced a concrete limitation: the composite formula as published cannot, by construction, allow the carbon-monoxide channel alone to cross the alert threshold, meaning the strong Class C results in Table 2 implicitly depend on a secondary CO-specific mechanism that is not yet made explicit in the formulation. Reporting this limitation candidly, rather than adjusting the historical results to obscure it, is intended to give downstream readers and reviewers an accurate picture of exactly what has and has not been validated.

Considered together, the simulation across 500 generated scenarios yields 97.3% detection accuracy, 99.4% precision, 98.4% recall, a 0.4% false-positive dispatch rate, and sub-1.5-second end-to-end alert latency on the warm path, at an estimated per-unit hardware cost of $120. These figures position the system as a cost-accessible intervention for a population at significant and growing risk, while the sensitivity and limitation analysis in Section 6 provides a more complete picture of where the present design is robust and where it still depends on assumptions outside the formally stated risk-scoring equation.

The future scope of this work follows directly from the gaps identified above. In the near term, the CO-channel limitation should be resolved with an explicit secondary threshold or a revised, re-validated weighting scheme, and the Azure Function cold-start penalty—shown in Figure 2 to dominate the latency budget—should be eliminated via a Premium-tier pre-provisioned deployment ahead of any field trial. In the medium term, physical prototype construction, instrumented-kitchen data collection, and IRB-approved studies with elderly-care participants are necessary to replace the procedurally generated sensor traces used throughout Section 6 with empirical ground truth, and to characterize failure modes—sensor drift, cross-sensitivity, connectivity loss—that a simulation cannot surface. Longer term, on-device machine learning inference, occupancy-aware risk weighting, multi-modal caregiver alerting, and a formal regulatory and patent pathway represent the natural extension of this architecture from a validated reference design toward a deployable product. Physical prototype validation and regulatory pathway work, together with closing the CO-channel gap, constitute the primary next steps toward real-world deployment.

## 9. ACKNOWLEDGMENTS

The author thanks the open-source community for libraries and tools used in the simulation and prototyping workflow.

## 10. REFERENCES

[1] National Fire Protection Association, “Home Fires Involving Cooking Equipment,” NFPA Research, Quincy, MA, USA, Tech. Rep., 2023. [Online]. Available: https://www.nfpa.org/education-and-research/research/nfpa-research/fire-statistical-reports/home-fires-involving-cooking-equipment

[2] M. Chan, D. Esteve, C. Escriba, and E. Campo, “A review of smart homes—Present state and future challenges,” Comput. Methods Programs Biomed., vol. 91, no. 1, pp. 55–81, Jul. 2008, doi: 10.1016/j.cmpb.2008.02.001.

[3] P. Rashidi and A. Mihailidis, “A survey on ambient-assisted living tools for older adults,” IEEE J. Biomed. Health Inform., vol. 17, no. 3, pp. 579–590, May 2013, doi: 10.1109/TITB.2012.2234129.

[4] J. Kim and S. Lee, “CookSafe: Non-intrusive load monitoring for kitchen appliance safety,” in Proc. IEEE Int. Conf. Consumer Electron. (ICCE), Las Vegas, NV, USA, 2019, pp. 1–4, doi: 10.1109/ICCE.2019.8661974.

[5] Y. Wang, H. Li, and Q. Zhang, “iStove: Vision-based flame and hazard detection for smart kitchens,” IEEE Sensors J., vol. 21, no. 8, pp. 10112–10121, Apr. 2021, doi: 10.1109/JSEN.2021.3056893.

[6] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, “Fog computing and its role in the Internet of Things,” in Proc. 1st Edition MCC Workshop Mobile Cloud Comput. (MCC), Helsinki, Finland, 2012, pp. 13–16, doi: 10.1145/2342509.2342513.

[7] Microsoft Corporation, “Azure IoT Reference Architecture,” Microsoft Docs, 2022. [Online]. Available: https://learn.microsoft.com/en-us/azure/architecture/reference-architectures/iot

[8] European Parliament and Council, “Regulation (EU) 2015/758 Concerning Type-Approval Requirements for the Deployment of the eCall In-Vehicle System,” Official Journal of the European Union, 2015.

[9] Occupational Safety and Health Administration, “Carbon Monoxide Poisoning,” U.S. Department of Labor, 2023. [Online]. Available: https://www.osha.gov/carbon-monoxide-poisoning