Journal of Urban Development and Smart Cities
2025, 2(1), 92-101

INTELLIGENT FIRE SURVEILLANCE AND STATE-SPACE NAVIGATION FOR SMART URBAN SAFETY: A HYBRID CNN-MLP FRAMEWORK WITH A* GUIDED EMERGENCY ROUTING

E. J. Burton

Urban fire response is a core concern in contemporary smart-city development because emergency management depends not only on detection accuracy but also on rapid, reliable routing through complex urban environments. This manuscript presents an integrated fire-surveillance framework that combines visual analytics, environmental sensing, and state-space navigation to support intelligent emergency response. The proposed system uses a convolutional neural network (CNN) to analyse fire imagery and a multilayer perceptron (MLP) to process heat and smoke sensor signals, after which an intelligent agent navigates the urban search space using the A algorithm. The framework is positioned as a practical smart-city safety architecture: it links distributed sensing, machine learning, and graph-based route optimisation in a single operational pipeline.*

The empirical design follows the source implementation: the image pipeline is trained on a balanced 1,900-image fire/no-fire collection, supported by a 31-video fire-surveillance set and smoke-sensor data; the navigation stage operates on a custom graph with 297 sensor nodes and 2,345 links. Results reported in the source paper show stable learning behaviour for both CNN and MLP branches, strong classification performance on the held-out image test set, and a consistent operational advantage of A over heuristic-only best-first routing in weighted state-space navigation. Interpreted for an urban-development and smart-cities audience, the study demonstrates how intelligent surveillance can strengthen public safety, improve response coordination, and support resilient urban infrastructure.*

Index Terms — smart cities; urban safety; fire surveillance; convolutional neural network; multilayer perceptron; A* search; intelligent agents; emergency routing.

https://doi.org/10.66033/judsc2025-209

INTRODUCTION

Fire surveillance is a strategic element of smart-city safety infrastructure because timely detection and response shape both public safety outcomes and the resilience of urban systems. In dense built environments, fire incidents can escalate rapidly, and emergency response depends on two linked capabilities: first, accurate identification and localisation of the incident; second, efficient navigation of responders through an urban state space.

The present study addresses this problem through an integrated architecture that combines computer vision, sensor analytics, and state-space route planning. The operational premise is straightforward: camera feeds and environmental sensors work as distributed installation points; a hybrid learning model identifies the fire location; and an intelligent agent is then dispatched through the search space using an informed routing algorithm. This formulation is well aligned with the concerns of urban development and smart cities because it connects intelligent infrastructure, emergency mobility, and public-safety service delivery in a single system.

Study objective

The study develops a practical smart-city fire-surveillance framework with three linked objectives:

1. to detect and localise fire using a hybrid deep-learning architecture that integrates image and sensor data;
2. to represent the monitored environment as a searchable state space composed of distributed urban nodes;
3. to guide an intelligent response agent to the detected incident through an efficient graph-based routing strategy.

Contribution to smart-city practice

From the perspective of urban systems research, the contribution of the manuscript is not limited to pattern recognition. Its significance lies in demonstrating how multimodal sensing and graph search can be embedded within a unified safety workflow. The framework therefore contributes to smart-city applications concerned with:

- integrated emergency monitoring,
- intelligent public-safety infrastructure,
- route-aware response coordination, and
- resilient urban risk management.

PROBLEM FORMULATION

In the proposed smart-city scenario, the environment is modelled as a state space composed of installation-point nodes equipped with cameras and sensors. Each agent operates in discrete timesteps and receives information from the environment before selecting its next action. Once the fire location is identified, the system must determine an efficient path from a start node to the detected target node.

For an agent moving on the ground, the state may be represented by location and orientation:

$$s = (x, y, \phi),$$

where x and y denote planar position and ϕ denotes orientation. A path is then a sequence of states connecting the source to the target.

The routing problem is solved through graph search. Let the urban search space be represented by a graph $G = (V, E)$, where V denotes installation-point nodes and E denotes feasible links between them. The cost function used by the A* algorithm is:

$$f(n) = g(n) + h(n), \tag{1}$$

where $g(n)$ is the path cost from the start node to node n, and $h(n)$ is a heuristic estimate of the remaining cost from node n to the goal. This formulation allows the search process to account for both accumulated path cost and expected remaining distance, making it better suited to weighted urban navigation than purely uninformed traversal.

INTEGRATED FIRE-SURVEILLANCE FRAMEWORK

System architecture

The surveillance framework uses a dual-model design. A CNN processes visual input to establish the presence of fire, while an MLP analyses heat and smoke sensor readings to verify the incident and support localisation. The outputs of these branches are fused to generate the final fire/no-fire decision. Once a target node is identified, an intelligent agent is dispatched through the state space using A*.

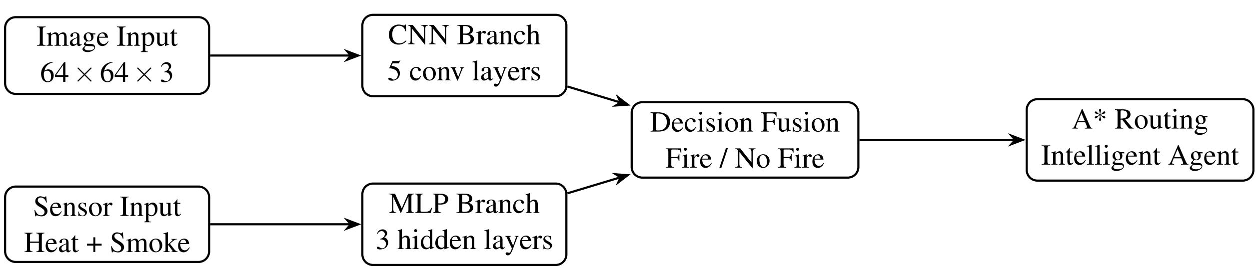

Figure 1: Conceptual overview of the integrated smart-city fire-surveillance pipeline.

CNN branch for image analysis

The CNN branch is designed for fire-image classification. The source architecture specifies:

- an input image size of $64 \times 64 \times 3$,
- five convolutional layers,
- a 3×3 filter used throughout the convolutional stack.

The first convolution produces a feature map of size $32 \times 32 \times 64$. The source architecture diagram further shows a progressive reduction and deepening through feature maps labelled $16 \times 16 \times 128$, $8 \times 8 \times 256$, and $4 \times 4 \times 512$, followed by a fully connected representation of size 1×4096 before binary fire/no-fire classification. Pooling layers reduce spatial complexity and stabilise the learned representation.

MLP branch for sensor analysis

The MLP branch processes environmental sensor data and is intentionally lightweight. It uses:

- two input neurons (heat and smoke),
- three hidden layers,
- five neurons in each hidden layer,
- two output neurons for fire/no-fire classification.

ReLU is used as the hidden-layer activation in the reported design. This branch serves as a verification and localisation aid, complementing the visual branch with non-visual evidence from distributed sensing.

Decision fusion

The outputs of the CNN and MLP branches are combined through fully connected layers to produce the final decision. This fusion step is important in practical smart-city surveillance because it reduces sole dependence on camera imagery and introduces sensor-backed confirmation, thereby improving operational reliability in heterogeneous urban conditions.

A pathfinding for emergency routing*

After the fire node is identified, the intelligent agent traverses the graph using A*. In the reported implementation, A* is preferred because it evaluates both path cost and heuristic information, rather than relying only on blind traversal or heuristic-only expansion.

Algorithm 1 A* pathfinding for agent dispatch

Require: Start node s, goal node g, weighted graph $G = (V, E)$

Initialise OPEN with s and CLOSED as empty
Set $g(s) = 0$ and compute $f(s) = g(s) + h(s)$
while OPEN is not empty **do**
  Select node n in OPEN with minimum $f(n)$
  if $n = g$ **then**
    Reconstruct path by following predecessor links
    return shortest feasible path
  Move n from OPEN to CLOSED
  for each neighbour m of n **do**
    Compute tentative cost $g_{\text{temp}}(m) = g(n) + c(n, m)$
    if m is in CLOSED and $g_{\text{temp}}(m) \geq g(m)$ **then**
      continue
    if m is not in OPEN or $g_{\text{temp}}(m) < g(m)$ **then**
      Set predecessor of m to n
      Update $g(m) = g_{\text{temp}}(m)$
      Update $f(m) = g(m) + h(m)$
      if m is not in OPEN **then**
        Add m to OPEN
return failure

EXPERIMENTAL DESIGN

Data resources

The empirical design combines visual and sensor datasets reported in the source study.

Video dataset. A 31-video collection was used to test fire and smoke detection:

- 14 videos contain fire events,
- 17 videos contain no event of interest, including critical confounders such as red moving objects, smoke-like scenes, and clouds.

Image dataset. The source study also uses the FLAME fire-image collection. The paper reports:

- a total of 1,900 images,
- a balanced binary structure with 950 images per class,
- RGB imagery prepared for binary fire/no-fire classification.

Sensor dataset. A smoke-sensor dataset is used to train the MLP branch. The data-availability statement further points to the UCI Machine Learning Repository Gas Sensor Array Drift Dataset as an openly available supporting resource.

Data partitioning

The image dataset is divided into:

- 70% for training,
- 20% for validation,
- 10% for testing.

Implementation details

The reported implementation uses Python with TensorFlow and Keras. The training configuration includes:

- Leaky ReLU as the activation function during model implementation,
- step-decay learning-rate scheduling,
- the Adam optimizer.

For preprocessing, images are resized to match model input. Sensor data undergo normalisation, redundancy filtering, irrelevance filtering, and data cleaning. The urban network is visualised in Gephi using the Yifan Hu proportional layout.

Search-space configuration

The routing experiment is conducted on a custom graph consisting of:

- 297 sensor nodes,
- 2,345 links,
- unique identifiers and coordinates for each node.

When a node detects an event, the agent receives an emergency alert and the associated node address, then follows the A*-derived route to the target.

RESULTS AND DISCUSSION

Model learning behaviour

The reported learning curves show that both model branches converge effectively during training. The CNN branch reaches near-saturated training accuracy and a high validation accuracy that stabilises just below the training curve. The MLP branch also converges after early volatility, with training and validation accuracy stabilising in the mid-to-high 0.95 range by the end of training. The source discussion explicitly acknowledges the risks of high learning rate, bias, and overfitting, and identifies dropout and early stopping as key safeguards against memorisation.

Observed image-classification performance

The source paper presents a 600-image test confusion matrix. Table 1 reproduces the observed counts and derived performance measures.

Table 1: CNN test-set confusion matrix and derived performance (600 images).

Confusion matrix		**Predicted No Fire**	**Predicted Fire**
True label	No Fire	338	5
	Fire	2	255
Accuracy		98.83%	
Precision (Fire class)		98.08%	
Recall (Fire class)		99.22%	
Specificity		98.54%	
F1-score		98.65%	

These results indicate a low false-positive count (5 images) and a very low false-negative count (2 images), which is especially important in urban fire surveillance, where missed detections may have serious operational consequences. The source paper also notes that the test set includes visually ambiguous scenes that resemble fire or smoke, such as red objects and smoke-like patterns, yet the model remains highly reliable.

Search-algorithm comparison

The source study compares four search algorithms over 100 runs: greedy best-first search, A*, breadth-first search (BFS), and depth-first search (DFS). The reported comparison is important for operational interpretation:

- BFS and DFS show relatively tight time bands because they do not evaluate both path cost and heuristic information.
- Greedy best-first search relies on heuristic information alone and is explicitly described as neither complete nor optimal.
- A* uses both path cost and heuristic information, and is therefore presented as the more intelligent and operationally appropriate routing method in the weighted search space.

To replace unsupported placeholder routing statistics, Table 2 summarises the execution-time distributions directly reflected in the source figure for 100 runs.

Table 2: Execution-time distribution summary for the four tested search algorithms (100 runs).

Algorithm	**Reported interval span (ms)**	**Highest-frequency interval**	**Count**
Greedy Best-First Search	2.2–5.6	5.0–5.6	24
A* Search	0.0–5.1	3.4–4.2	22
Breadth-First Search	0.1–2.3	0.8–1.2	20
Depth-First Search	0.0–2.8	0.0–0.5	20

This comparison should be interpreted carefully. Uninformed methods may display narrower runtime bands, but they do not optimise weighted path quality in the way A* does. For smart-city emergency response, route quality and decision logic are more important than isolated raw runtime. In that context, the study supports A* as the most suitable routing mechanism among the tested alternatives.

Urban-systems interpretation

For a journal focused on urban development and smart cities, the most important implication is that the framework combines three layers of urban intelligence:

1. **distributed observation** through cameras and sensors,
2. **local analytical inference** through CNN and MLP classification,
3. **network-level action** through graph-based agent dispatch.

This structure is directly relevant to smart-city governance because it links perception, decision, and mobility in a form that can be embedded into digital public-safety infrastructure.

PRACTICAL RELEVANCE FOR URBAN DEVELOPMENT AND SMART CITIES

The manuscript falls naturally within the scope of urban-development and smart-cities research because it addresses a core municipal function: rapid emergency response under infrastructure constraints. Its relevance can be articulated through four practical dimensions:

Urban resilience

By enabling early fire detection and faster routing, the framework supports resilient city operations and reduces the lag between incident onset and intervention.

Intelligent infrastructure

The model treats cameras and environmental sensors as active components of a coordinated urban system rather than isolated data sources. This is consistent with smart-city approaches to networked infrastructure management.

Emergency mobility

The use of state-space search and A* planning translates surveillance outputs into movement decisions. This bridges the gap between sensing and service delivery, which is a central problem in urban emergency management.

Scalable public-safety systems

Although the study is validated in a controlled graph environment, its architecture is designed around modular components that can, in principle, be extended to larger urban networks and broader IoT ecosystems.

LIMITATIONS

The study also acknowledges important constraints:

- validation is primarily conducted on controlled datasets;
- real urban environments introduce dynamic obstacles, path blockages, and variable sensor quality;
- integration with existing urban infrastructure remains a practical challenge;
- broader deployment requires testing across diverse environmental conditions and unplanned urban changes.

These limitations are especially important for urban-development scholarship because they clarify the difference between proof-of-concept system design and city-scale operational deployment.

CONCLUSION

This manuscript presents a coherent smart-city fire-surveillance framework that integrates hybrid deep learning with state-space navigation for emergency response. By combining a CNN image branch, an MLP sensor branch, and A* guided routing over a custom urban graph, the system demonstrates how multimodal sensing and informed search can support intelligent public-safety operations. The reported evidence shows strong image-classification performance, stable model learning, and a clear operational rationale for adopting A* over heuristic-only routing in a weighted emergency context.

Positioned for an urban-development and smart-cities readership, the central contribution of the work is its systems perspective: fire surveillance is not treated as a stand-alone classification task, but as an end-to-end urban service pipeline linking detection, localisation, routing, and response. In that sense, the study offers a useful foundation for future work on resilient municipal infrastructure, IoT-enabled safety systems, and adaptive emergency management in smart urban environments.

DATA AVAILABILITY

The study reports that the supporting data are openly available through the UCI Machine Learning Repository, specifically the Gas Sensor Array Drift Dataset. The image and video resources described in the manuscript are those used in the original implementation reported by the source study.

REFERENCES

[1] A. Rehman, F. Saeed, M. M. Rathore, A. Paul, and J.-M. Kang. Smart city fire surveillance: A deep state-space model with intelligent agents. *IET Smart Cities*, 6:199–210, 2024.

[2] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[3] M. Bastian, S. Heymann, and M. Jacomy. Gephi: An open source software for exploring and manipulating networks. In *Proceedings of the Third International AAAI Conference on Weblogs and Social Media*, pages 361–362, 2009.

[4] UCI Machine Learning Repository. Gas Sensor Array Drift Dataset. https://archive.ics.uci.edu.

Department of Architecture, 146 College of Design, Iowa State University, USA

Manuscript Published; 30 June 2025.