

FIREScore: A FRAMEWORK FOR INCIDENT RISK EVALUATION, SIMULATION, COVERAGE OPTIMIZATION AND RELOCATION EXPERIMENTS

Guido Legemaate¹, Joep van den Bogaert², Rob van der Mei³, and Sandjai Bhulai⁴

¹Div. of Information Services, Fire Dept. Amsterdam-Amstelland, Amsterdam, NETHERLANDS

²Jheronimus Academy of Data Science, 's-Hertogenbosch, NETHERLANDS

³Dept. of Stochastics, Centrum Wiskunde & Informatica, Amsterdam, NETHERLANDS

⁴Dept. of Mathematics, Vrije Universiteit Amsterdam, Amsterdam, NETHERLANDS

ABSTRACT

This paper introduces fireSCore, an open source framework for incident risk evaluation, simulation, coverage optimization, and relocation experiments. As a digital twin of operational fire department logistics, its visualization frontend provides a live view of current coverage for the most common fire department units. Manually changing a unit status allows for a view into future coverage as it triggers an immediate recalculation of prognosed response times and coverage using the Open Source Routing Machine. The backend provides the controller and model, and implements various algorithms, e.g. a relocation algorithm that optimizes coverage during major incidents. The data broker handles communication with data sources and provides data for the front- and backend. An optional simulator adds an environment in which various scenarios, models and algorithms can be tested and aims to drive current and future organizational developments within the Dutch national fire service.

1 INTRODUCTION

Fire departments worldwide face operational challenges of providing timely emergency responses with often limited resources, especially personnel, across expanding urban areas. In addition strategic and tactical decisions regarding station locations, vehicle allocations, and dispatching policies significantly impact response times and operational efficiency. However, these decisions are often complex, requiring evaluation of multiple interacting factors and numerous possible configurations. Simulation methods have proven valuable for analyzing emergency service systems, allowing planners and decision-makers to evaluate potential configurations without disrupting actual operations.

The framework presented in this paper enables, on an operational level, the live visualization of current coverage by various types of fire units. The ability to manipulate vehicle status supports decisions by the dispatcher to allow units to be temporary unavailable, e.g. to leave their service area for training purposes. Other components enable testing and modeling of vehicle dispatching decisions and resource allocation with the help of a simulator, which simulates emergency incidents based on historical data and statistical forecasting.

The remainder of this paper is organized as follows: Section 2 reviews related work in relevant areas of emergency service and coverage optimization. Section 3 describes the organization of the Dutch fire services, including the response times that have to be met and a short primer on current developments. Section 4 describes the various components of the framework architecture. Section 5 describes the various statistical, behavioral, and expert-driven methods used to validate the simulator. Finally, Section 6 ends with a conclusion.

2 LITERATURE REVIEW

Fire departments, and emergency services in general, have been a likely candidate for employing operations research (OR) techniques and simulation methods to optimize their logistics, including station locations, resource allocation, routing, and workload balancing to - indeed - make every second count. Mathematical optimization models are used to find the best fire station locations and to (re)allocate resources, often evaluated with simulation. Recent years have seen a surge in advanced approaches – including discrete-event simulations, agent-based models, and digital twins – that incorporate real-time data and complex system dynamics.

2.1 Fire Station Location Models

Optimizing the locations of fire stations is a classic problem in emergency services, often with the goal to minimize the response times to incidents or maximize coverage of demand locations. Traditional OR models like the Maximal Coverage Location Problem (MCLP) and p-median formulations have been widely used. A comprehensive survey by Aleisa (2018) lists numerous approaches for fire station location, including fuzzy multi-objective optimization, maximal coverage models, GIS-based analysis, genetic algorithms, ant colony optimization, Tabu search, and simulated annealing.

More recently, researchers have extended location models to be more dynamic, data-driven, and multi-faceted. Multi-objective and stochastic models are increasingly common. Tao et al. (2022) introduced a backup coverage optimization model to ensure that high-risk areas have multiple stations covering them. In their case study of Wuhan, China, the model increased the proportion of high-fire-risk zones covered by more than one station – from about 38.5% (with the current 85 stations) to over 50% when optimally placing 95 stations – thereby enhancing redundancy in coverage. These multi-covering approaches address reliability issues as stations might be busy or otherwise unavailable when needed, by providing overlap in coverage. Similarly, Ming et al. (2022) proposed a distributionally robust optimization model for station planning under uncertainty. Their model jointly optimizes station locations, the number of fire trucks, and demand-area assignments by minimizing the worst-case expected total cost, considering uncertainties in incident demand and travel times.

A European contribution to the station location problem is a Dutch example by van den Berg et al. (2017). They developed a location-allocation model for, and in cooperation with, the Amsterdam-Amstelland fire department that not only decides station locations, but also extended the model by allocating different types of fire units to each station. A case study revealed that relocating just 3 of the city's 19 fire stations could cut the fraction of late arrivals by over 50%, without adding any new stations.

2.2 Emergency Vehicle Routing and Response Time Optimization

Once stations are in place, the next operational challenge is to facilitate emergency vehicle routing to incident locations as quickly and safely as possible. Route optimization in this context deals with finding the fastest paths to minimize travel time. In practice, fire departments typically rely on shortest-path algorithms with travel time predictions; however, research shows that more sophisticated strategies and real-time controls can yield improvements. A comprehensive review by Hao et al. (2024) classified the literature into three areas: travel time prediction, routing optimization, and traffic priority control. Key findings of this review were that accurate travel-time prediction (e.g. using historical data or machine learning) and active priority measures (like smart traffic signals) can significantly aid routing.

In the Netherlands, Usanov et al. (2020) studied the problem of dispatching multiple fire trucks when an incident occurs, recognizing that standard protocol (sending the nearest unit) might not always be optimal when driving times are uncertain or when multiple units are required. They formulated the fire truck dispatching problem as a Markov Decision Process and computed optimal policies for a case study with the Amsterdam-Amstelland fire and safety region. One key scenario is when two units must respond to the same incident, which is a common operational policy for building fires in the dense city center of the

Dutch capital. Results showed that by intelligently choosing which station's units to dispatch, sometimes not the being the absolute nearest, the fraction of late arrivals could be reduced by 20% on average, and over 50% in certain high-traffic scenarios, compared to the nearest-unit policy.

2.3 Resource Allocation and Workload Balancing

Effective fire service operations require not only good station locations and routing, but also optimal allocation of resources and balanced workloads. Resource allocation in this context can mean determining how many fire or specialized units to station at each fire department, how to assign crews, and even how to schedule or relocate units dynamically to cover gaps. Workload balancing refers to ensuring that no station or unit is overwhelmed with calls while others are underutilized.

Several optimization models incorporate these aspects. The study by van den Berg et al. (2017) mentioned earlier is an example of a location-allocation model: it simultaneously decides on station locations and the distribution of different vehicle types across those stations. By including multiple vehicle types with different response time targets, the model handles the allocation of specialized resources (for example, heavy rescue or aerial units might have different coverage requirements than standard fire units, like pumpers). Moreover, by accounting for crews with mixes of full-time (professional) and part-time (volunteer) firefighters, it implicitly balances workload and staffing – part-time availability can vary by time of day, so the model ensures that each station's deployment meets the coverage targets with the available personnel. This is particularly pertinent in Europe, where many fire services (especially in smaller cities and rural areas) rely on part-time firefighters and must carefully balance their deployment. For instance, some Dutch regions found that offering 24/7 coverage with part-time crews was challenging, leading to experiments with different crew sizes and station timings to meet an 8-minute response time standard (Koppenjan et al. 2019).

Next to static allocation, there is an increasing interest in dynamic resource management: temporarily reallocating or redistributing coverage when certain units are busy, a practice common in ambulance services known as 'move-up' or dynamic deployment. In the fire service dynamic relocation is less common, in the Dutch context likely due to different workload characteristics when compared to ambulance services. However, large-scale incidents might leave areas uncovered, requiring neighboring units to reposition. Some recent studies have begun to simulate such scenarios. For instance, Usanov et al. (2019) studied dynamic fire truck relocations during major incidents, modeling the strategic repositioning of fire units to maintain balanced coverage across the Amsterdam-Amstelland region when substantial resources are committed to a single major incident for an extended period of time. Their analysis demonstrates the potential effectiveness of dynamically adjusting vehicle positions to ensure improved coverage and reduced response times in affected areas.

2.4 Forecasting

In addition to simulation and optimization models, accurate forecasting of emergency incidents supports to proactively manage fire service operations. It helps to predict future demand, informing decisions on resource allocation, station placement, and staffing, further enhancing response effectiveness and readiness. Given the geographical and regional nature of incident prediction, this section exclusively focuses on forecasting studies conducted in the Dutch context, for two different incident types. In the Amsterdam-Amstelland region a forecasting study aimed to predict the number of incidents handled by fire stations, emphasizing the impact of severe weather conditions on operational services. It distinguishes between small and major incidents, with major events modeled as an inhomogeneous Poisson process due to their sporadic nature. For small incidents, several modeling approaches are compared, including a Linear Model (LM), a Generalized Linear Model (GLM) that incorporates interactions among weather variables, a Random Forest (RF), and an ensemble averaging method. Weather conditions, particularly wind and rainfall, are shown to significantly influence incident frequency, while temperature exhibits non-linear effects and

visibility contributes marginally. The ensemble model, combining GLM and RF outputs, demonstrates the best predictive performance (Legemaate et al. 2021). In the more rural region of Twente a study was undertaken to forecast the number of chimney fires. In a two-step approach, random forests with conditional permutation importance to non-parametrically select key environmental and building-related variables were used. Subsequently, a nested Poisson point process model that separates spatial and temporal risk components was developed. Validation through second-order statistics and residual analysis confirms the model's effectiveness in representing the spatio-temporal patterns of chimney fire incidents (Lu et al. 2023).

2.5 Synthesis

The literature reviewed in this section highlights the integral roles of station location optimization models, emergency vehicle routing, resource allocation, and forecasting in enhancing fire department operations. Each research area contributes distinct yet complementary perspectives: station location optimization models provide robust foundational coverage, advanced routing techniques ensure swift response, dynamic resource allocation and workload balancing models guarantee operational flexibility, and accurate forecasting methodologies allow proactive resource planning. The functional and architectural design of fireScore facilitates in the integration of these operational insights, optimization methods, and data-driven forecasting approaches, providing a unified decision-support framework.

3 THE DUTCH FIRE SERVICE

The Dutch fire service maintains a network of approximately 950 stations staffed by about 23,000 firefighters – roughly 80% of whom are volunteers. This reliance on part-time crews, especially in rural areas, means that many firefighters respond from home or work when alerted, incurring a slightly longer turn out time, but it enables cost-effective staffing. Each region develops coverage plans to meet target response times for its jurisdiction; for high-priority incidents, typical response times are on the order of 8 minutes (IFV 2021).

To support strategic planning and ensure rapid response, the Dutch fire services started to employ simulation tools and optimization models. Researchers and practitioners have applied location allocation models and discrete event simulations to optimize station locations, vehicle deployment, and crew allocation, aiming to minimize response times.

3.1 Safety Regions

The fire services in the Netherlands are organized at a national-regional hierarchy established by law. In 2010 the Safety Regions Act (SRA) transformed municipal fire brigades into 25 regional fire departments, called safety regions, with the goal to coordinate fire and rescue services across multiple municipalities (Ministry of Justice and Security 2010). National oversight and policy guidance are provided by the Ministry of Justice and Security and bodies like the Council of Fire Commanders, while each safety region manages its own operations including a regional dispatch center for emergency calls. Emergency calls are handled by these joint control rooms, which given the incident classification dispatch the nearest available fire units of the proper type. An important task is to ensure that coverage is maintained for subsequent calls.

3.2 Response Times

With the introduction of the SRA and its accompanying decree (SRD), legally mandated fire department response time norms have been formally established and made binding. The SRD classifies structures into four risk-based categories by occupancy type, with corresponding maximum response times of 5, 6, 8, and 10 minutes from highest to lowest risk, respectively as shown in Table 1.

Table 1: Maximum response times (for fire incidents) by building type.

| Building Type Category | Maximum Response Time |
|---|----------------------------------|
| Enclosed shopping centers, residences above shops, detention facilities | 5 minutes |
| Apartment buildings, housing for persons with reduced self-reliance | 6 minutes |
| Other residential or retail types, healthcare, educational, or lodging facilities | 8 minutes |
| Offices, industrial sites, sports/assembly venues, other functions | 10 minutes |
| <i>Justified deviation (with regional rationale)</i> | <i>18 minutes (absolute max)</i> |

Under the SRA, each safety region must develop a coverage plan that defines the applicable response time norms for every area in the region. This plan must adhere to the national standards set by the SRD; any deviation (i.e. a longer response time) is only permitted under strict justification by the regional authorities, and in no case may a locally adopted norm exceed an absolute maximum of 18 minutes.

In addition, a comprehensive registration and monitoring of actual response times is mandated to verify that the fire brigade’s performance meets these norms in practice; notably, the statutory response time norms apply only to high-priority incidents requiring an urgent emergency response.

To operationalize these response time standards, the coverage plan is translated into Station Deployment Priority Tables (SDPTs). These partition the region into predefined demand or dispatch zones. For every zone, one or more SDPTs exist. These tables define the ordered sequence in which fire units are sent based on the type of incident and the required vehicle type. The configuration of SDPTs allows for differentiation by context, such as daytime hours versus nights and weekends. SDPTs are created by vehicle type, such as standard fire engines (pumpers), water rescue units, aerial apparatus, technical rescue vehicles, and other specialized units. Certain SDPTs enable strategic coordination, such as dual-directional approaches or mutual aid protocols across regions. More recently, dynamic SDPTs have been added which incorporate real-time vehicle availability and Global Positioning Systems (GPS) to enable more adaptive dispatching. Incidents are geographically mapped to one of these zones by the emergency dispatch system and subsequently the required fire units are sent using the order in the SDPTs.

3.3 Current Developments

The Dutch fire service is undergoing significant transformations to enhance its effectiveness and adaptability in response to evolving demographic needs and complex emergencies. One development is the shift towards *area-based response times*, moving away from rigid, object-specific standards to a more flexible, risk-oriented approach. Furthermore it introduces new performance measures next to (adapted) response times, such as operational capacity, readiness and workload that help fire departments communicate their performance and resource needs more clearly (Brandweer Nederland 2022).

Also, the fire service is embracing the concept of the *adaptive fire service*, aiming to develop an organization that is agile and responsive to rapid societal and environmental changes. This involves diversifying functions and tasks, implementing flexible staffing and deployment models, possible introducing additional vehicle types. Another aspect is enhancing coordination through advanced operational centers, likely on the subject of supporting (national) operational capacity during major incidents. The goal is to develop a resilient organization capable of effectively managing both current and emerging challenges given changing demographics (NIPV 2022).

4 COMPONENTS

The framework consists of 4 main components as shown in Figure 1. A frontend, acting as a digital twin geared towards fire department operations, delivers a real-time visualization of the operational environment, displaying incidents, vehicle location and status, and expected coverage. The framework incorporates a data broker to retrieve and process relevant operational data. A controller model backend manages internal logic and supports the implementation of various algorithms.

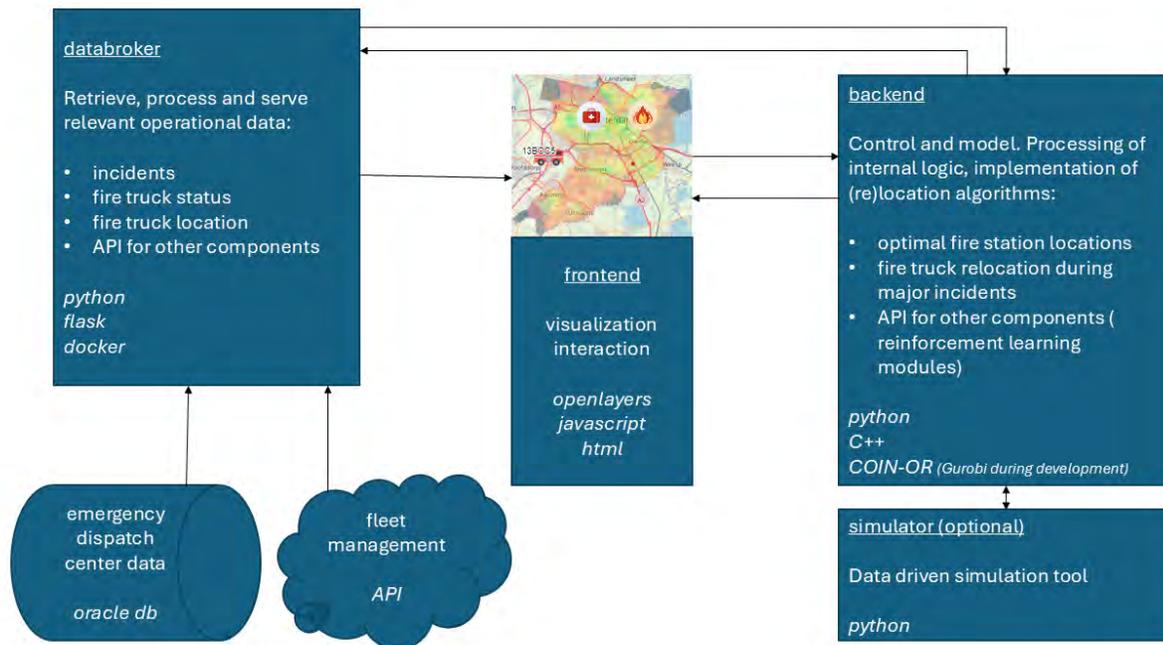


Figure 1: System architecture and components (emergency dispatch center data and fleet management are necessary data sources).

4.1 Visualization Frontend

The frontend integrates geospatial mapping capabilities with real-time incident data representation. Using information derived from the data broker, such as current (availability) status of fire units, each demand location is colorized based on the SDPT. Color coding indicates the level of coverage across the demand locations as can be seen in Figure 2. Colors and thresholds can be adjusted on the client side, but by default reflect the maximum response time categories as can be seen in Table 1. It utilizes a client-server architecture constructed using HTML5, CSS3, and JavaScript, with the OpenLayers library (version 10.4.0) serving as the primary geospatial rendering engine. This enables the representation of incident data on a dynamic map interface utilizing the Dutch national coordinate system (EPSG:28992). Incident data is visually encoded through vector features overlaid on a base map. Each incident is represented as a point feature with an icon representing the incident type. Associated metadata, including geographic coordinates, priority classification, and alarm text is available upon request. The system dynamically manages these features, clearing and regenerating the visualization layer upon receipt of updated data, ensuring accurate representation of the current operational coverage in both a static (SDPT) and dynamic (GPS) way.

The interface effectively implements coordinate system transformation through the Proj4js library, enabling accurate positioning of incidents on the standardized Dutch reference system while maintaining interoperability with global web mapping standards. The system implements a real-time data visualization approach through WebSocket connections, allowing for immediate reflection of backend data changes without page refreshes. This asynchronous communication method significantly enhances user experience by providing low-latency updates to critical emergency management information.

4.2 Controller Model Backend

The backend includes a controller that processes most of the internal logic. From preprocessing raw incident data into structured data frames with calculated incident rates used for the (relocation) algorithms,

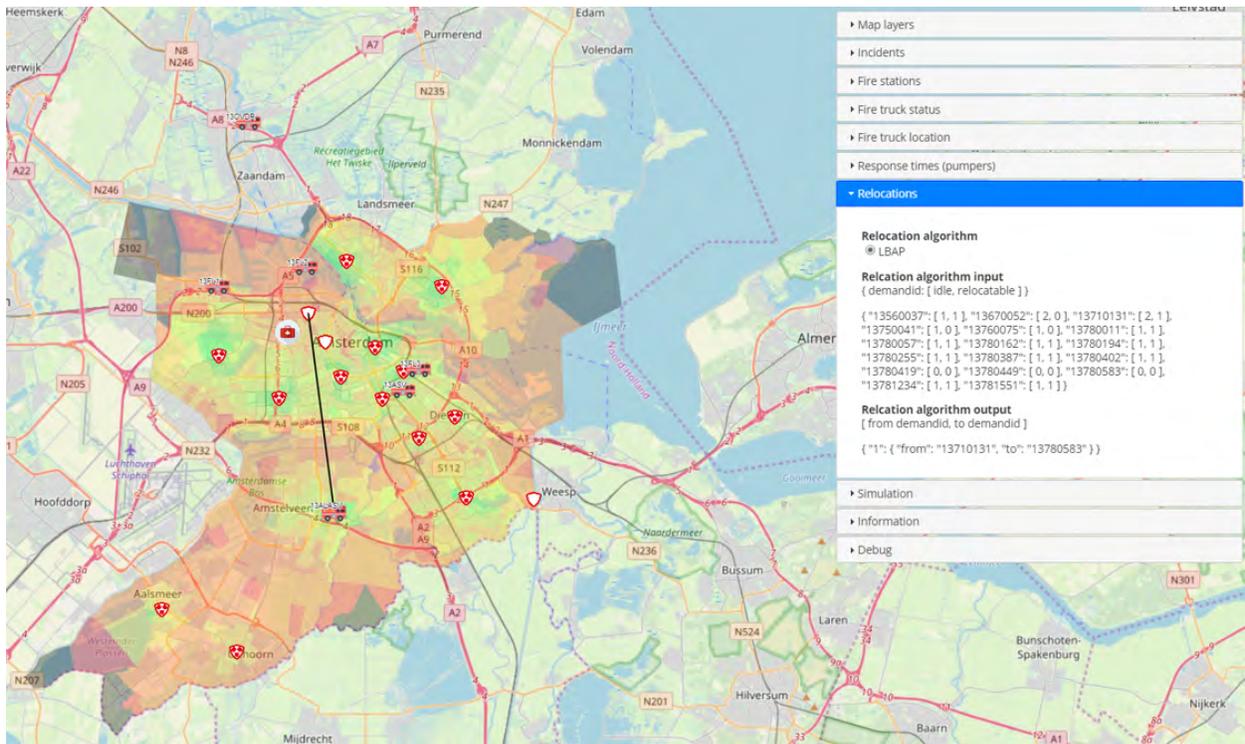


Figure 2: Visualization frontend with response times (colors), fire stations and units (pumpers), incidents and relocation proposal.

or travel time matrices to support fast and accurate simulations to the actual implementation of some of these algorithms.

One example of such an algorithm currently implemented is a fire truck relocation model. This algorithm can be triggered when there is a lower than usual amount of pumpers available and remaining pumpers need to be repositioned, e.g. during major incidents. The algorithm consists of two sequential steps. First a Maximum Coverage Relocation Problem (MCRP) is solved, followed by a Linear Bottleneck Assignment Problem (LBAP).

The MCRP balances three main goals: (i) maximizing geographic coverage across the service region, (ii) limiting the number of truck movements, and (iii) ensuring fair distribution of resources. The model achieves this by taking into account the current availability of trucks, the spatial distribution of incident demand based on historical data, and a parameter that captures the dispatcher's wish to be 'the best boy in class'. This willingness parameter allows for balancing operational efficiency with performance improvements.

A distinctive feature of the MCRP formulation is its fairness constraint, which is expressed through the concept of *response neighborhoods*. These are sets of demand locations defined by proximity to the nearest n fire stations. The model ensures that each neighborhood is covered by at least one idle truck, preserving equity in emergency service access even during periods of resource shortages. The model is formulated as an integer program that specifies, for each idle truck, whether and where it should be relocated, subject to constraints on station capacities and coverage obligations.

Once the MCRP identifies which relocations should occur, the LBAP is solved to determine how to efficiently assign specific trucks to specific destination stations. The LBAP minimizes the maximum relocation time among all truck movements, thus accelerating regaining full coverage. This model treats the relocation task as a one-to-one matching between origin and destination stations, ensuring that no truck is delayed excessively during redeployment.

The sequential use of MCRP and LBAP ensures that both the strategic coverage goals and operational efficiency considerations are addressed. MCRP generates the optimal relocation plan from a coverage perspective, while LBAP ensures that the execution of that plan is achieved with minimal latency. This combined approach is particularly effective in dynamic emergency environments, where maintaining responsive and equitable service coverage is critical under constrained resources (Usanov et al. 2019). The algorithm is implemented using the excellent PuLP library, a linear and mixed integer programming modeler written in Python (Mitchell et al. 2011). As a solver the fire department, being a public service, gracefully uses the open source COIN-OR linear programming solver written in C++ (Lougee-Heimer 2003).

In the visualization the optimal solution, in regard to the willingness, is communicated in one or more lines that are drawn from the source to the destination station(s) as show in Figure 2. In addition the solutions, if possible, given a slightly lower or higher willingness are made available so the dispatcher can manually evaluate if another solution might be beneficial given the whole environment.

4.3 Data Broker

The data broker manages and processes data streams using a microservices architecture. A key objective is to offload high-frequency data retrieval tasks from operational systems, such as the emergency dispatch center, reducing load and improving overall system responsiveness for its clients. It integrates multiple components; core components being a scheduler, a database, and a message broker. The system is containerized using Docker. The scheduler is responsible for executing periodic tasks defined as jobs. It uses the Python APScheduler library for scheduling and managing jobs. Jobs are defined in Python modules. Each job consists of a *Getter* (data retrieval logic), an optional *Emitter* (data broadcasting logic), and an optional *Persister* (data storage logic). The scheduler communicates with other components via Redis, an in-memory key-value pair database, used as a distributed cache and message broker. If set, messages can be made persistent using MongoDB, a document-oriented NoSQL database.

4.4 Simulator

The data-driven simulator supports decision-making in the context of fire service logistics. It integrates several key modules, each responsible for simulating and analyzing a different aspect of incident response, with a primary focus on dispatch, travel times, vehicle availability, and incident forecasting. The simulator operates as a modular and extensible platform for replicating realistic emergency incident response dynamics, its core functionalities include:

- **Incident Generation:** Incidents are sampled stochastically using a Poisson process, with time-dependent intensity rates derived from historical data. These rates are forecasted using Facebook's *Prophet* model, allowing for dynamic and time-sensitive incident sampling.
- **Incident Characterization:** For each sampled incident, key attributes are drawn probabilistically: type of incident, spatial location, priority level, building function, and the required vehicle mix, although the latter can also be fixed based on the incident type. These are based on empirical probability distributions fitted to historical fire department data.
- **Response Time Modeling:** Response times are decomposed into dispatch time, turnout time, and travel time. The dispatch and turnout components are modeled using parametric distributions (e.g., lognormal and gamma), where turnout times are additionally fitted per station and incident type. Travel time is estimated using Open Source Routing Machine (OSRM) travel durations, adjusted with a multiplicative gamma-distributed noise factor to reflect stochastic travel delays (Luxen and Vetter 2011).

Dispatch times are sampled uniquely per incident, rather than per individual deployment. The simulator generates a single dispatch delay value for each simulated incident, representing the delay from incident notification to the alarm. After this dispatch delay, multiple vehicle deployments corresponding to the incident can subsequently occur. Thus, while the turnout and travel times may

differ per deployment (vehicle-specific and station-specific), the initial dispatch delay is sampled once per incident, reflecting the centralized dispatching process realistically and consistently across all deployments responding to that particular incident.

- **Vehicle Dispatch and Movement:** Vehicles are assigned to incidents using a *ShortestDurationDispatcher*, which selects the fire station that minimizes expected travel time. The simulator updates vehicle status and location over time, enforcing availability constraints and queuing behavior in the case of limited resources. Alternative assignment methods can easily be implemented by adding a new dispatch rule.
- **Spatial and Functional Context:** Demand locations are initialized with distributions over building functions, enabling the simulator to reflect spatial heterogeneity in risk and incident characteristics.
- **Simulation Execution:** The simulator can run for a predefined number of incidents or over specific time periods. Each incident triggers the full response cycle, including resource assignment, delay sampling, and performance logging.

The simulation engine supports experimentation with operational policies (e.g., relocation strategies), infrastructure changes (e.g., station openings/closures), or forecasting scenarios. It provides a realistic and flexible foundation for evaluating emergency response performance under varying conditions.

4.5 Practical Usage

The simulator is used as a comprehensive decision-support tool that enables the fire department to evaluate and improve performance across strategic and tactical levels, leading to operational improvements. By modeling the spatio-temporal distribution of incidents and simulating a realistic emergency response processes, it can provide quantitative insights into the consequences of policy and resource allocation decisions.

At a strategic level, the simulator is used to assess long-term planning scenarios. One application is the evaluation of candidate locations for new fire stations. Through simulation of incident occurrence and response dynamics, planners can identify alternative locations that yield improvements in coverage and/or response time performance. Furthermore, the simulator can support investment decisions by helping to quantify the expected benefits of adding new vehicles to the fleet, including their location. These capabilities allow planners to iteratively develop coverage plans aligned with changing regional risk profiles, legal response time requirements and adapting to changing societal needs.

At a tactical level, the simulator is used to explore alternative resource configurations within an existing infrastructure. This includes the assessment of different vehicle-to-station allocation strategies to enhance spatial coverage and/or response times. It facilitates the analysis of personnel scheduling policies, particularly in mixed systems where both full-time and part-time crews are utilized. Scenarios can be modeled to evaluate the implications of part-time staffing during specific hours or days, or to examine the effects of temporary or permanent station closures. It enables the evaluation of dynamic vehicle relocation strategies, allowing comparison of algorithms that reposition vehicles in response to coverage deficits. Additionally, the simulator can be used to analyze different dispatch rules, assessing how variations in deployment heuristics affect overall performance. These evaluations helps fire departments to adapt their operations to financial constraints and changing risk profiles while maintaining service quality.

5 VALIDATION

Validation of the simulator involves a combination of unit testing of the code, expected behavior (including statistical comparisons), and ongoing expert judgment and review from domain experts within the fire department. A basic comparison of historical response time distributions based on real-world data and response time distribution outputs from the simulator can be seen in respectively Figures 3 and 4.

Data used for this visualization was in both cases based on 10 years of high-priority incidents for the incident type *building fire* and the time used was taken from the first pumper arriving on-scene, consistent with current regulations. The outliers, mainly on the upper ends, can be attributed to those stations assisting

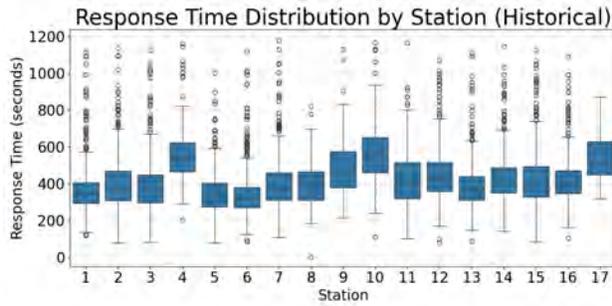


Figure 3: Historical response time distribution.

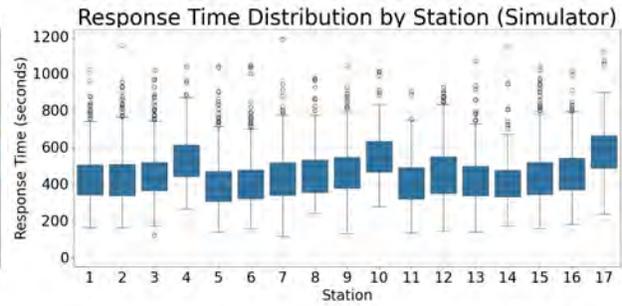


Figure 4: Simulated response time distribution.

fires outside of their own service area, needing longer travel times resulting in longer response times. Response times higher than 1200 seconds (20 minutes) were removed from the historical dataset as analysis showed that these were mainly due to data registration errors. Stations 1, 2, 3, 6, 7, 11, 13, 14, 15 and 16 are 24/7 staffed with full-time firefighters. Station 5 is also staffed by full-time firefighters, but this station closes between 11 pm and 7 am. Station 12 is a genuinely mixed station with a hybrid of full-time and part-time firefighters. When out for a call, part-time firefighters are called to the station to restaff remaining fire units. Stations 4, 9, and 10 are also part-time stations by designation, but during office hours aimed to be staffed with part-time firefighters who can work remotely from the fire station. For stations 4 and 10, this policy has only been introduced recently. Stations 8 and 17 are completely staffed with part-time firefighters only. Every fire station has at least one pumper unit on call.

Further statistical validation was done by analyzing the three distinct components that make up the response time: dispatch time, turnout time, and travel time. The simulated dispatch and turnout time distributions align closely with historical data, as confirmed by Kolmogorov-Smirnov tests yielding p-values greater than 0.05, indicating no statistically significant differences, as shown in Figures 5 and 6. Simulated travel times for certain station–demand-location routes exceed those in historical data, yielding marginally longer journeys than expected. This discrepancy arises either due to limited occurrences in real-world data, reducing representativeness, or because the simulation employs OSRM-based travel time estimates using a standard car profile, which does not account for differences in speed between regular traffic and fire trucks across varied local conditions. Further research will utilize GPS-based travel times to develop a custom *emergency services* profile for OSRM, enhancing travel time accuracy.

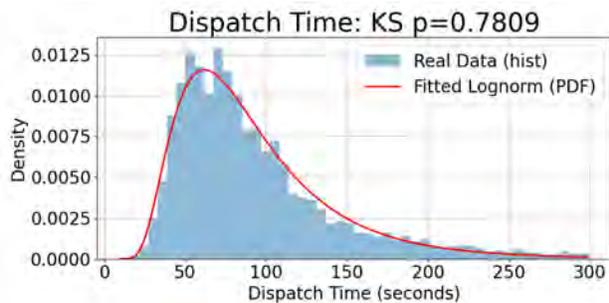


Figure 5: Dispatch time distribution

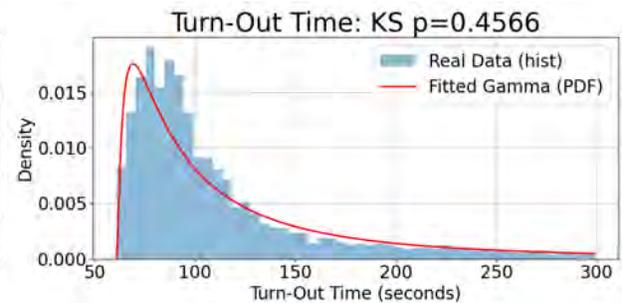


Figure 6: Turnout time distribution

Validation by domain experts included six practical experiments assessing the simulator’s response to operational changes, such as modifying vehicle availability, adjusting station locations and operating hours, adding or removing stations, activating a backup protocol, and updating forecasts. In 14 out of 15 scenarios, the simulator behavior aligned with expert expectations; only one station-relocation case showed a minor deviation.

For example, removing a pumper fire truck from a high-demand city-center station led to a significant increase in overall response times: the mean rose from 08:08 to 08:19 minutes, and the 95th percentile from 13:55 to 14:12 minutes. Within the affected service area, delays for high-priority incidents also worsened, with mean delays increasing by 1 minute and 4 seconds, and the 95th percentile rising from 46 seconds to 1 minute and 41 seconds. All effects were statistically significant ($p < 0.001$).

6 CONCLUSION

The fireScore framework offers a comprehensive, integrative approach that involves incident risk evaluation, simulation, coverage optimization, and relocation experiments within a single platform. By integrating real operational data with a data-driven simulation engine, it - much like a digital twin - provides a powerful decision-support tool capable of real-time coverage visualization and scenario-based analysis of emergency response strategies. The modular framework supports the Dutch fire service's current performance metrics and can easily be adapted to facilitate the ongoing shifts towards area-based, risk-oriented response strategies and adaptive service models, introducing the evaluation of new performance metrics (e.g., operational capacity, readiness and workload).

Beyond this, the added benefit of a simulator also serves a diagnostic role by revealing structural weaknesses in the current emergency response system and those proposed by the current developments within the Dutch fire service, aiding further improvements in the emergency response system as a whole.

ACKNOWLEDGMENTS

We thank all students who contributed to this and related projects, particularly Wenjie Hu, for assistance in preparing the various graphs presented in this publication.

REFERENCES

- Aleisa, E. 2018. "The Fire Station Location Problem: A Literature Survey". *International Journal of Emergency Management* 14:291–302 <https://doi.org/10.1504/IJEM.2018.094239>.
- Brandweer Nederland 2022. "Gebiedsgerichte Opkomsttijden en Uniforme Dekkingsplannen (Area Based Response Times and Uniformal Coverage Plans)". Dutch Fire Services. <https://www.brandweernederland.nl/onderwerpen/gebiedsgerichte-opkomsttijden-en-dekkingsplannen>, accessed 14th February 2025.
- Hao, Z., Y. Wang, and X. Yang. 2024. "Every Second Counts: A Comprehensive Review of Route Optimization and Priority Control for Urban Emergency Vehicles". *Sustainability* 16(7):2917 <https://doi.org/10.3390/su16072917>.
- IFV 2021. "Brandweerkerndata 2020 (Statistics on Dutch Fire Services 2020)". Instituut Fysieke Veiligheid (Institute for Physical Safety). <https://nipv.nl/wp-content/uploads/2022/03/20211109-IFV-Brandweerkerndata-2020.pdf>.
- Koppenjan, J., A. van Sluis, F.-B. van der Meer, B. Kuipers, and K. van Paridon. 2019. "Changing the Rotterdam-Rijnmond Regional Fire Service: A Multi-Dimensional Approach for Analysing Strategies of Change". *International Journal of Emergency Services* 9(2):123–141 <https://doi.org/10.1108/IJES-04-2018-0023>.
- Legemaate, G. A. G., J. de Deijn, S. Bhulai, and R. D. van der Mei. 2021. "Severe Weather-Based Fire Department Incident Forecasting". *International Journal on Advances in Intelligent Systems* 14:14–23.
- Lougee-Heimer, R. 2003. "The Common Optimization INterface for Operations Research: Promoting open-source software in the operations research community". *IBM Journal of Research and Development* 47(1):57–66 <https://doi.org/10.1147/rd.471.0057>.
- Lu, C., M.-C. van Lieshout, M. de Graaf, and P. Visscher. 2023. "Data-Driven Chimney Fire Risk Prediction using Machine Learning and Point Process Tools". *Annals of Applied Statistics* 17(4):3088–3111 <https://doi.org/10.1214/23-AOAS1752>.
- Luxen, D., and C. Vetter. 2011. "Real-time routing with OpenStreetMap data". In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. November 1st-4th, Chicago, Illinois, 513–516. <https://doi.org/10.1145/2093973.2094062>.
- Ming, J., J.-P. P. Richard, R. Qin, and J. Zhu. 2022. "Distributionally Robust Optimization for Fire Station Location under Uncertainties". *Scientific Reports* 12(1):5394 <https://doi.org/10.1038/s41598-022-08887-6>.
- Ministry of Justice and Security 2010. "Wet Veiligheidsregio's (Safety Regions Act)". Ministerie van Justitie en Veiligheid. <https://wetten.overheid.nl/BWBR0027466/2025-02-12>, accessed 12th February 2025.
- Mitchell, S., M. O'Sullivan, and I. Dunning. 2011. "PuLP: A Linear Programming Toolkit for Python". Technical report, Department of Engineering Science, The University of Auckland.

- NIPV 2022. “Toekomstverkenning Brandweer 2022-2030 (Future Outlooks Fire Services 2022-2030)”. Nederlands Instituut Publieke Veiligheid (Netherlands Institute for Public Safety). <https://nipv.nl/wp-content/uploads/2022/10/20221012-NIPV-Toekomstverkenning-Brandweer-2022-2030.pdf>.
- Tao, L., Y. Cui, Y. Xu, Z. Chen, H. Guo, B. Huang *et al.* 2022. “Location Optimization of Urban Fire Stations Considering the Backup Coverage”. *International Journal of Environmental Research and Public Health* 20(1):1–18 <https://doi.org/10.3390/ijerph20010627>.
- Usanov, D., G. A. G. Legemaate, P. M. van de Ven, and R. D. van der Mei. 2019. “Fire Truck Relocation During Major Incidents”. *Naval Research Logistics* 66(2):105–122 <https://doi.org/10.1002/nav.21831>.
- Usanov, D., P. M. van de Ven, and R. D. van der Mei. 2020. “Dispatching Fire Trucks under Stochastic Driving Times”. *Computers & Operations Research* 114:104829 <https://doi.org/10.1016/j.cor.2019.104829>.
- van den Berg, P., G. A. G. Legemaate, and R. D. van der Mei. 2017. “Increasing the Responsiveness of Firefighter Services by Relocating Bases in Amsterdam”. *Interfaces* 47(4):352–361 <https://doi.org/10.1287/inte.2017.0897>.

AUTHOR BIOGRAPHIES

GUIDO LEGEMAATE is a researcher, coordinator of the information services division and on-scene information manager at the Amsterdam-Amstelland Fire Department. He collaborates with various scholars and students to improve operational fire service processes through scientific insights. Over the years, they have developed various optimizations and algorithms. He holds an M.Sc. in Health Sciences and is currently pursuing a Ph.D. at CWI as an external PhD candidate. His research focuses on applying data science to improve emergency response services. He co-authored a study on increasing the responsiveness of firefighter services by relocating base stations in Amsterdam, published in the journal *Interfaces*. Another contribution is his work on fire truck relocation during major incidents, published in *Naval Research Logistics*. Through his interdisciplinary approach, combining practical experience with academic research, he continues to contribute to the enhancement of emergency response services in the Netherlands. His email address is guido.legemaate@veiligheidsregioaa.nl and his website is <https://www.linkedin.com/in/gagleemaate>.

JOEP VAN DEN BOGAERT is an industrial engineer and data scientist, holding an MSc in Industrial Engineering and an Engineering Doctorate (EngD). He developed the original version of *fdsim*, the simulation engine for modeling fire department incidents and responses in the Netherlands, as part of a collaboration between the Fire Department Amsterdam-Amstelland and the Jheronimus Academy of Data Science. His expertise encompasses full-stack development and data science, contributing to advancements in emergency response strategies. Currently he is an engineer at and the co-founder of *Rewize*, a company focused on innovating e-commerce supply chains. His email address is joepvdbogaert@gmail.com and his website is <https://www.linkedin.com/in/joepvandenbogaert>.

ROB VAN DER MEI is a Full Professor of Applied Mathematics at the Vrije Universiteit Amsterdam, and senior researcher and Manager of Research and Strategy at Centrum Wiskunde & Informatica (CWI). He holds M.Sc. degrees in Mathematics and Econometrics from the Vrije Universiteit Amsterdam and earned a Ph.D. in queueing theory from the University of Tilburg in 1995. His research focuses on Stochastic Operations Research, bridging theory and application in areas such as emergency response systems, predictive policing, traffic management, ICT scalability, service and freight logistics, and data analytics. He has established extensive industry and government collaborations, leading numerous public-private partnerships. He initiated and leads the *REPRO* project on proactive ambulance planning. He has also chaired the European COST project *ACROSS* and currently leads the *DOLCE VITA* project on optimizing elderly care in the Netherlands. In 2021, he and Prof. Sandjai Bhulai received the prestigious *Huibregtsen Award* for research combining scientific excellence with societal impact. He has co-authored approximately 250 publications in peer-reviewed journals and conferences. His e-mail address is mei@cw.nl and his website is <https://www.cwi.nl/en/people/rob-van-der-mei>.

SANDJAI BHULAI is a Professor of Business Analytics at Vrije Universiteit Amsterdam. He studied Mathematics and Business Mathematics & Informatics before earning a Ph.D. in Markov decision processes for complex systems control. He co-founded the Amsterdam Center for Business Analytics and the postgraduate Business Analytics/Data Science program. His research focuses on decision-making under uncertainty, optimization, and data science, with applications in HR analytics, predictive modeling, and dynamic pricing. His current projects include social media analytics, AI-driven workforce management, and planning and scheduling in complex systems. He has co-authored approximately 100 publications in peer-reviewed journals and conferences. A recipient of multiple teaching awards, he is also a frequent keynote speaker and panelist in both academia and industry. His email address is s.bhulai@vu.nl and his website is <https://research.vu.nl/en/persons/sandjai-bhulai>.