Master thesis 2022

TOOL PLACEMENT AND SMART NOTIFICATION IN THE CIVIL VOLUNTEER NETWORK

Follow-up research for the Fire Department Amsterdam-Amstelland on the future civilian network

Amsterdam, November 8, 2022

Vrije Universiteit Amsterdam
Fire Department Amsterdam-Amstelland
Christiaan Smit
Master's program Business Analytics

cj.smit@student.vu.nl

Supervised by:
First supervisor: Rob van der Mei, Vrije Universiteit Amsterdam
Company supervisor: Guido Legemaate, Fire Department Amsterdam-Amstelland
Second reader: Mathisca de Gunst, Vrije Universiteit Amsterdam

1 Management summary

The main topic of this thesis is the future civil volunteer network for the Fire Department Amsterdam-Amstelland (FDAA). A civil volunteer network can make a crucial contribution to the traditional fire department by arriving before the fire department and performing simple tasks to ease the work for the arriving fire fighters. However, setting up this network leads to many important choices in the design of the system. The goal of this thesis is to identify some of these choices and to measure their impact.

To investigate the impact of certain policies, a realistic theoretical model is built in which several experiments are performed. These experiments are meant to represent these different policies and also attempt to give a best- and worst-case outlook on the results. The main policies and options that were modeled and tested are:

- **Smart notification**. This topic deals with the question: "*How many* and *which* volunteers should receive a notification to respond to an incident?" Options for this policy are *area notification*, which sends out notifications in an area of 100m around the incident location, and *k-closest notification*, which selects the k closest volunteers to an incident.

- **Differentiation between volunteer types**. This thesis makes a distinction between three different types of volunteers: *rapid responders*, *tool carriers*, and *specialists*. Rapid responders directly move to an incident, whereas tool carriers first pick up tools at a predetermined tool location. Specialists are a subset of the volunteer population with a specific kind of expertise. They can assist the fire department once that expertise is necessary for an incident.

- **Differentiation between incident types**. One of the main differences between the fire department and other emergency services is the diversity of the incidents responded to. This is taken into account in the model by restricting certain tool types to certain incident types.

The performance of these policies was tested on simulated data and deployments. On average, this resulted in a time gain of 4 to 7 minutes for rapid responders compared to the fire department. This same average time gain amounted to $1\frac{1}{2}$ to 3 minutes for tool carriers. Considering average response times from the fire department are around 10 minutes, these numbers are very impressive. Specialists had a mean response time of 10 to 16 minutes. The best results were achieved in the more populated central area of the Amsterdam-Amstelland safety region.

Out of the two smart notification policies, k-closest notification was clearly superior. It resulted in better time gains, while also decreasing the average number of responding volunteers. The latter is an important metric, because when too many people respond to the same incident, this can lead to less motivated volunteers and perhaps lower acceptance rates.

Note that the results and options in the model do not necessarily prescribe the choices that the FDAA should make in the process of designing the volunteer network. Instead, this research serves as an opportunity to investigate certain choices in a controlled environment. Moreover, the code in which the model was built is also one of the end products to the FDAA, with which future experiments can be performed once more information on the design of the future network is available.

Contents

2 Introduction

2.1 Fire Department Amsterdam-Amstelland

The Fire Department Amsterdam-Amstelland (from here on also referred to as FDAA) is a branch of the national Dutch Fire Department operating in the safety region Amsterdam-Amstelland, which includes the city of Amsterdam and the municipalities of Aalsmeer, Amstelveen, Diemen, Uithoorn and Ouder-Amstel, as can be seen in Figure 1. Weesp was also added to the municipality of Amsterdam in March 2022, but is left out of this research due to lack of historical data. Although region Amsterdam-Amstelland is a relatively small safety region, it is one of the regions with the highest population. Still, there are many areas in the safety region with very low population density, such as the port area in the northwest and the less urbanized areas in the south. The FDAA responds to all sorts of incidents and in fact fires are only a small part of those incidents (approximately a quarter of the urgent responses). Other types include cardiac arrests, traffic incidents and storm and water damages. To accommodate those different incidents, the FDAA has many different vehicles distributed over the 18 base stations. In recent years, the department has become more data-driven, especially with the help of the team Information Management (IM) and the collection of data which has already lead to many developments in the research area, in close collaboration with entities such as the Vrije Universiteit Amsterdam and Centrum Wiskunde & Informatica (CWI). However, the organization as a whole is also very much concentrated on the practical side of business, which sometimes makes it difficult to actually implement improvement of processes. So, to make this research a successful one, it is essential to clearly explain the practical value of the models and conclusions.

Figure 1: *Map of Amsterdam-Amstelland safety region*

2.2 Civil volunteering

The main topic of this thesis is the civil volunteering program that the FDAA wants to introduce. It is important to note the difference between these volunteers and the on-call volunteers, who are essentially regular firefighters with the only notable difference that they are called from home instead of the station. The civil volunteers are not employed by the fire department, but have an app installed which can alert them to respond to a nearby incident. In the rest of this thesis, the term 'volunteer' should be read as 'civil volunteer', unless stated otherwise. Part of the long-term policy plan for 2030 includes the 'ferocious goal' of 100,000 civil volunteers that the FDAA plans to recruit to their program.

The main task of civil volunteers is to respond to incidents in order to provide assistance to the firefighters or to perform simple tasks beforehand so that the fire department can focus on dealing with the incident itself. However, there are more benefits of the civil volunteering network than just responding to incidents. For example, the volunteers can also be tasked to help neighbors by giving advice on fire safety and checking their smoke detectors. The network can also be used as an information channel, where the fire department can spread tips on fire safety and fire prevention in society through personal contacts. This is especially desirable if people can be contacted through the volunteer network that are otherwise hard to reach via traditional channels. As these effects are hard to model, they are out of scope for this research, but it is still important to remember that such a program can cause more societal impact than it seems initially.

The civil volunteer program can also be seen from a different perspective than that of response time: that of mobility. The city center has by far the highest density of incidents and is also home of the two busiest stations. However, typical fire trucks are generally not made for its narrow streets. They damage the embankment of the canals by their weight and can get blocked up by traffic or unexpected obstacles. An obvious solution is to deploy faster and more lightweight motorcycle units, especially when no heavy equipment is needed for the incident. However, volunteers fulfill a very similar purpose by already being very close to the incident to begin with and responding by foot or by bike. And instead of bringing tools with them, as in a fire truck, these tools can be distributed over the region to be picked up by volunteers. This thesis aims to model this as realistically as possible.

2.3 Positioning in organization

This thesis serves as a *proof of concept* for the future civil volunteer program. This means that all the experiments and scenarios tested are possible within the future volunteer program, but are not necessarily required to be implemented. Instead, this thesis aims to provide numerical calculations to scenarios that could possibly correspond to the way the network will be implemented in practice. In these scenarios, several options are given for possible policies and their results in terms of response time are discussed. After this, the FDAA can decide based on the results whether this policy is indeed valuable to the program or not. As such, it is the responsibility of the FDAA to develop the program and also to decide on how the network will function in practice. The latter has been and will be discussed in this thesis, but only in general terms without many details, as this is not the scope of this research. Selecting candidate options for policies had been done in discussion with experts within the FDAA by identifying which challenges the fire department is currently facing and will face in the future. If this is a challenge that the civil volunteer program might be able to solve, it is fitted into the model and added to the experiments.

2.4 Problem statement

This thesis aims to extend on the previously mentioned research by factoring in tools that can be picked up by the volunteers. It also aims to model the notifications and volunteer responses more realistically, which in turn leads to many challenges regarding notification strategies. This leads to the main research question:

What is the impact of policies such as smart notification and tool placement in the civil volunteering program on the response time of the volunteers?

This question can be split into the following subquestions:

1. What is the best way to model volunteers picking up tools?
2. How can volunteers be notified in a smart way so that they are only sent out when they are really needed?
3. What are the benefits and costs of smart notification?
4. What is the optimal placement of tools over the service area?
5. How can we distinguish between different types of tools and how will this impact the response time of tool carriers?

2.5 Thesis outline

In Section 3, relevant literature will be discussed, including research performed for the FDAA as well as general research in the topic of civil volunteering. Section 4 covers the datasets that are used in the experiments and provides some interesting data exploration. Section 5 explains the outline of the model used in this research and aims to justify the choices made in creating this model. Section 6 gives an outline on the experiments performed within the model and how these should help in answering the research questions that were stated before. In Section 7 the results of the experiments are discussed and concluded in Section 8, where the results will also be compared to the research questions. Lastly, Section 9 will attempt to form the implications of the conclusions as well as further research on this topic or similar topics.

3 Literature research

In this section, all literature relevant to this research will be discussed. Section 3.1 and 3.2 will cover most of the relevant previous researches at the FDAA, as well as some interesting researches in other fire departments. The main topic of Section 3.3 is research in civil volunteering. Because nearly all of the examples of existing civil volunteering programs are in the medical field, this will be the main focus.

3.1 Schönberger research

This thesis serves as the follow-up from *Schönberger (2021)* [1]. In this thesis for the FDAA, response times were calculated for the volunteer network under different scenarios. These scenarios were created by changing certain parameters, such as the number of volunteers, the speed of volunteers, or the pre-trip delay. Another main topic of this research is the spatial distribution of the volunteers over the safety region. This distribution was either proportional to the number of inhabitants or optimized to efficiently cover as many incidents as possible. The exact current implementation of this is discussed in Section 5.3. The main conclusion to this research was that the volunteer network seems to be a valuable addition to the current deployments when looking at response times. Moreover, the goal of 100,000 civil volunteers does not need to be reached to obtain good coverage. For example, a network with 50,000 volunteers that go by bike instead of by foot achieves similar gains in response time.

3.2 Fire department research

The fire department has been the subject of many previous optimization researches. These researches are not only interesting from a technical point of view, but also give insight in how the fire department works and which factors are important to take into account for modelling. One interesting topic is incident forecasting, which is obviously an important issue when working with incident data. One of the earliest examples of this is for the New York fire department in the early 1970s, where incident rates were predicted as a function of location, time, method of reporting, and weather [2]. A similar research has been done quite recently for the FDAA as well, where also different types of incidents were predicted with different measures (mostly weather factors) [3]. Another interesting observation is that New Year's Eve is by far the busiest day regarding incidents and is therefore treated as an outlier. This means that caution is required when taking this data into account, although this might become less of an issue with the stricter fireworks regulations that are recently implemented in the Netherlands.

Other research, by *van den Bogaert (2019)* [4] has resulted in the implementation of a simulator, specifically designed for simulating incidents and movements of fire trucks in the Amsterdam-Amstelland safety region and making use of a deep reinforcement learning algorithm. This simulator is also used in this research using the most recent incident and deployment data.

Another important field of optimization within the fire department focuses on the placement of fire stations and staff on a strategic level. Specifically within the FDAA, *van den Berg et al. (2017)* [5] constructed a model where the benefit of moving base stations could be calculated. By changing only a few base stations, coverage could be increased by around 5%, which resulted in a reduction of late arrivals by approximately 50%. *Usanov et al. (2019)* [6] focused on the relocation of fire trucks instead, where they can be moved to a different station when many trucks are busy due to a major incident. This is already common practice, but this has never been done methodically.

3.3 Medical volunteering research

Although civil volunteering is a newly emerging field for the fire department, in medical care there exist similar programs designed to respond to out-of-hospital cardiac arrests, for example GoodSAM [7] in the UK and PulsePoint [8] in the US. This field is especially interesting, because it tries to tackle similar problems as the fire department. In fact, many of the calls that the FDAA answers to are cardiac arrests. Also the topic of tool placement comes back in the form of AED (automated external defibrillator) placement. For example *Chan et al. (2016)* [9] considered different scenario's with multiple responders, single responders, and whether or not they are able to find the closest AED. *Sun et al. (2016)* [10]

adjusted the current coverage in Toronto by temporal availability (for example buildings closing at night) and found a significantly better distribution of AED's when taking into account this restriction. *Van den Berg et al. (2021)* [11] showed that the spatial distribution of volunteers could be modelled as a Poisson point process and they optimized the distribution accordingly.

The Netherlands also has a well-developed volunteer program for cardiac arrests, called HartslagNu [12]. This system currently has roughly 250,000 members distributed over the whole country and the program already has significantly increased survival rates. It can alert volunteers either by app (based on location) or SMS (based on home or work address) and it also includes possible pickups of AED's. HartslagNu is partnered with many researches to improve alerting. One example is Slaa (2020) [13], where a model was created using HartslagNu data and from that model, new alerting methods were found which in the end resulted in a 13% increase in survival probability. These methods included:

- Increasing the alert radius, but limiting the number of volunteers alerted.
- Sending the volunteer to the AED that minimizes the total walking distance, instead of the volunteer closest to an AED.
- Only cancelling alerts after all tasks are filled (instead of when enough people have accepted the alert).

Most of these changes are directly implemented into the notification methods used in this research.

3.4 Main contributions

This thesis aims to contribute and expand in a meaningful way to the previously named researches. These main contributions are:

- Making a distinction between multiple types of volunteers, by adding tool carriers and specialists as new roles to the currently known volunteer role. The concept of tool carriers is similar to that of people picking up AED's, but implementing this in the context of the fire department leads to new, unexplored, challenges.
- Differentiation between multiple tool types, and optimal placement of tool locations, which is not something that is heavily researched for AED's.
- Putting more focus is on new and smarter methods of notification. Currently, this subject was only found in Slaa (2020) during the literature research.

4 Data

In this section, the main data sources and their content are discussed. Although 4 datasets were provided (incidents, deployments, station locations and vehicle allocation), only the data containing the incidents and deployments will be considered. The locations of the stations are quite trivial and therefore not suitable for deep data exploration. Vehicle allocation is an interesting subject, but out of scope for this project. The incident and deployment data is taken from the FDAA data warehouse and cover a period from January 6 2008 to the current date (end of October 2022). The rest of this section will present the data cleaning and exploration on both datasets, as well as other data obtained for this research.

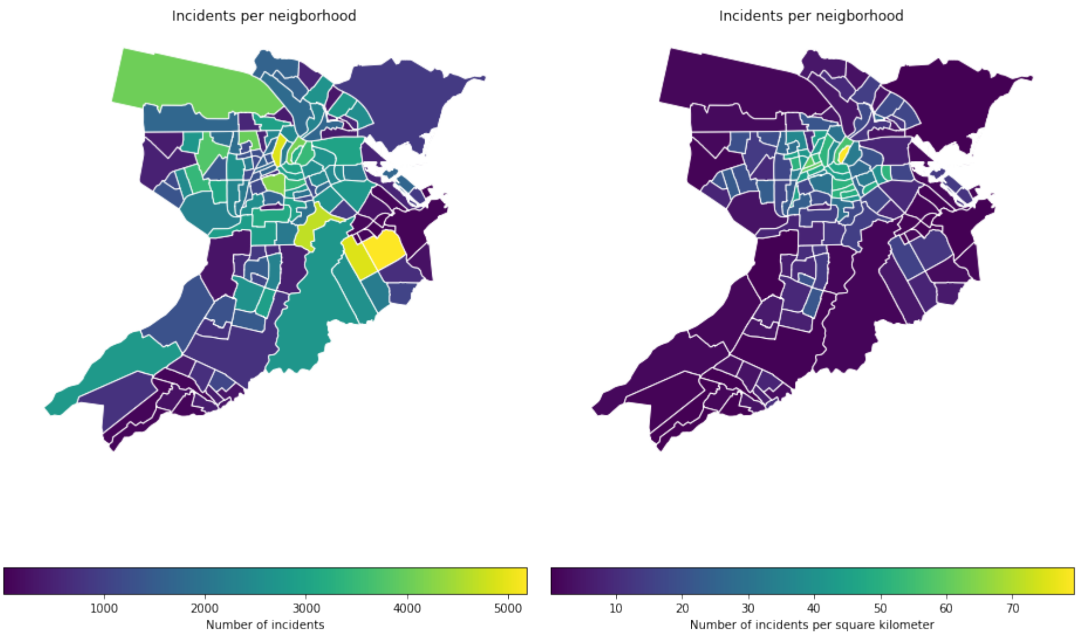

(a) *Counts* **(b)** *By area*

Figure 2: *Number of incidents by neighborhood: total counts vs per km*^{2}

4.1 Incident data

Figure 2a shows the number of incidents per neighborhood within the recorded period, and Figure 2b shows the same counts but divided by the area of the neighborhood. The map clearly shows a high density of incidents close to the city center. This is to be expected, as there is also the highest population density and economic activity. This is partially compensated by the fact that there are two stations close to the center, namely Nico and Hendrik, but still those stations are much busier compared to the stations located more towards the edge of the region.

Next to the frequency of the incidents, it is also important to investigate the different types of incidents. Figure 3 shows the different types that occur within the data and how often they occur. Note that for every type there exist more detailed description levels, but those do not give a general overview and would result in too many subdivisions of the data. Fire incidents and emergency service, the latter for example including people getting stuck in elevators or locked out of their home, are the most frequently occurring types of incidents. Other frequently occurring categories are healthcare, mainly cardiac arrests (besides and ambulance, fire and police department are also always called), and automatic alarm systems. Environment mainly consists of pollution or weather-related incidents.

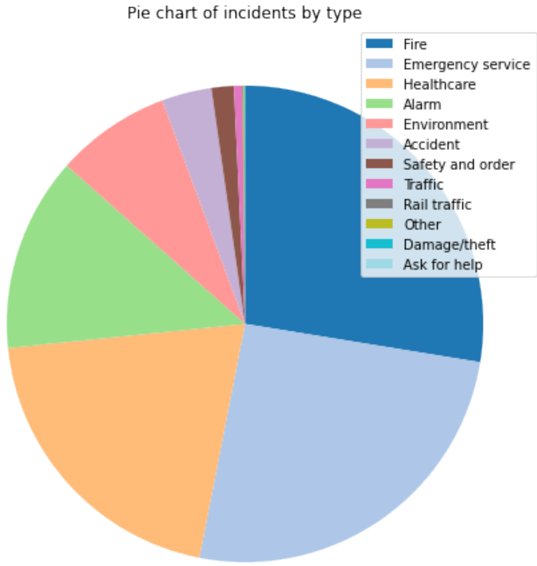

Figure 3: *Distribution of different incident types*

4.2 Deployment data

One of the most important aspects regarding to the research is the *response time of the deployments*. This response time is defined as the time between the moment when the call arrives at the dispatch center and the moment the fire department arrives at the incident. All urgent incidents have a target response time that is dependent on the type of building of the incident and the incident type. The percentage of responses where this target is met is one of the metrics used in this research and most similar researches. However, there are some problems with calculating the response times and the response time targets. First of all, there are cases where the building type is not available in the data. This might be because the data is simply missing, or because there is no building related to the fire and therefore there can be no target attributed. For these cases, a target of 8 minutes was given, because this is by far the most common target and most of the cases where the type was missing it seemed to involve just residential buildings. In the data, some of the response times were missing too. In those cases, the data was deleted as in those cases the time of the call was unavailable. In other cases, the response time was unrealistically high, which was almost always caused by an unusually high process time of the call. As these calls contribute to approximately 10% of the data, it is important to handle them well. After deliberation with an expert, it seems that these anomalies are likely caused by one of four (similar) causes:

1. The fire department is called after another emergency service. For example: the police are sent to investigate a report of smoke. If they indeed find something, the fire department is then called, but the time from the first call is taken as the response time.

2. Sometimes the emergency dispatch center leaves a registration open to collect extra data about the incident. It could happen that the incident turns out to require more assistance, in which case an extra unit is sent. This extra unit will then have a excessively long response time. This variant is quite rare compared to other variants.

3. Someone is called for scheduled ambulance transportation, but a ladder truck is needed to transport the person from a higher floor. The starting time of the call will then be the moment the transportation is scheduled. However, these incidents should already be filtered out as they are not urgent.

4. There has been a migration of the dispatch center to join it together with the police and ambulance. It is possible that this has resulted in some faulty data regarding response times.

Issues 1, 3 and 4 are solved by using the call time of when it is accepted by the fire department dispatcher instead of the general dispatcher. This should generally make a difference of only a few seconds, but not if any of these issues have occurred. For fairness, this is only done on calls that did not reach their target already. This resolves most of the issues with very high response times. The cases where the response time was still over 1 hour were then deleted from the data.

After this data cleaning, the total percentage of targets met is 80.8%. Figure 4a shows the development of the fraction of met target over the years, starting from the beginning of the data. Figure 4b shows the distribution of response times. Clearly, there are still some outliers, but there seem to be only very few meaning that they have no significant effect on the on-time rate. Most responses seem to be around 400 seconds and a large portion of the incidents are responded to within 10 minutes. This is already impressive, but this number can likely be outperformed by civil volunteers. Note that from Figure 4b, it is not possible to directly infer the percentage of on-time responses, as the targets can differ by incident.

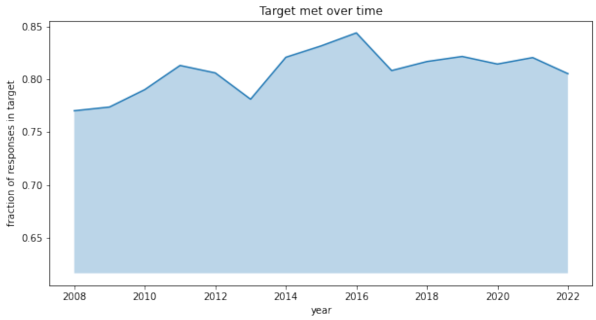

(a) *Distribution of different incident types*

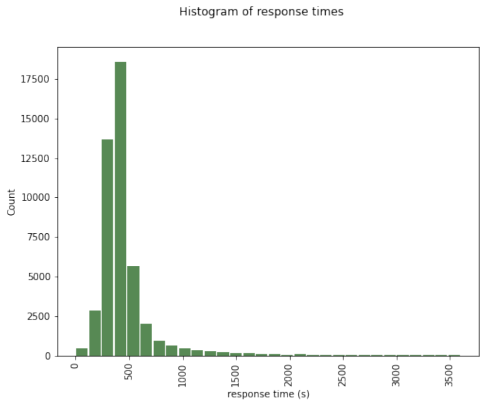

(b) *Histogram of response times*

Figure 4: *Plots of response times*

4.3 Other data sources

Alongside the data provided by the FDAA, some other datasets were used to be combined with the incident and deployment data. These dataset mainly include geospatial data of the safety region, as well as census data containing information about population numbers, etc. This data is publicly available and obtained via sources such as CWI (Centrum Wiskunde & Informatica), or similar internal data available at the FDAA.

4.4 Simulated data

The main data source on which the model was tested is created by a simulator which was a product from the research by *van den Bogaert (2019)* [4]. The logic behind this is that this results in enough data to train and test on, and also that future trends in the data are already accounted for. As there are still some outliers regarding response time in the simulated data, it was decided to remove the 1%-quantile of entries with the highest response time to give a more realistic view on the results. Unfortunately, the exact coordinates of the incident location are not simulated. For this reason, the coordinate was calculated by taking a random coordinate within the neighborhood of the incident. Since the incident rate will be assumed to be uniform over a neighborhood, this does not have any further implications. As for the response time target, this is set to 8 minutes for all incidents. This makes it easier to analyze the results and is not far from reality, since the vast majority of incidents have an 8 minute target time anyway.

4.5 Confidentiality considerations

Many companies are hesitant to publish their data in either raw or aggregated form, as well as researches that use this data in any meaningful way. However, as the FDAA is a governmental organization, typically relying on publicly available data, this is not a problem for this research and censoring is not necessary. In fact, the FDAA strives to be as transparent as possible to the public, as they are indirectly funded by them through taxes.

Another important consideration is privacy. A significant part of the data consists of incidents that people might not be comfortable with to share publicly. However, the data only contains basic location information and therefore difficult to relate to specific persons or households. Also, the findings presented in this thesis will be only neighborhood-based and can therefore not be traced back to a single address. Most of the results are also based on simulated data, which of course contains no private information.

5 Model

In this section, the main model of the research is described. In creating this model many choices are made in order to stay close enough to reality, while also making the model simple enough to perform experiments in. These choices and assumptions with their limitations and their influence on the outcome of the research are also discussed in this section. The model is programmed in Python and all features are part of multiple documented packages that together form one of the end products of this research.

In Section 5.1, the arrival of incidents is investigated. Section 5.2 explains how the different types of incidents are distinguished and what implications this has on for example tool placement. Section 5.3 discusses how the volunteer distribution is calculated and how one can correctly sample from this distribution. In Section 5.4, the different types of volunteers are stated, as well as the reasoning behind the choices for these types. Section 5.5 explains the choices behind the tool placement and provides the formal mathematical optimization models. Section 5.6 shows how volunteer responses and notifications are modeled.

5.1 Arrival of incidents

In general, the arrival of incidents is assumed to follow a Poisson process. This is a process that naturally occurs when events occur independently of each other. This co-independence of incidents is therefore one of the assumptions of the model, but a very natural one. These incidents arrive with rate λ_i^t, where i corresponds to the incident location and t to the time. In this research, every neighborhood (Dutch: buurt) was taken as a different incident location. This assumes that all locations within this neighborhood have the same arrival rate. As the neighborhoods are relatively small, and most neighborhoods consist of similar buildings, this assumption seems to be valid. These incident rates are calculated based on historical data and for most problems, such as tool placement or distributing the volunteers, their time dependency is not important as the problems have no dimension of time: it is assumed that tools cannot move during the day and it is also assumed that the volunteer distribution stays constant over time. The first is a valid assumption, as locations of tools are set locations that are not easily movable. The latter assumption could prove to be wrong, as volunteers probably reside more at residential areas during the nights and weekends, whereas they are often at work during weekdays. This could be an extension for further research, but for now this distribution is assumed to be constant.

5.2 Incident types

In comparison to other emergency services, the fire department responds to a much larger variety of incident types. This means that there is more practical value in distinguishing between different types of incidents. The categorization of these types is the same as the one explored in Section 4.1. In order to preserve flexibility, the model allows the FDAA to focus on a specific set of types only. This could be for example the 5 most prevalent incident types, but also a different predetermined selection. One important aspect here is that the types that are not accounted for are entirely ignored. So when calculating for example a distribution of tools for the set of types: $\{Fire, Alarm, Accident\}$, then the model assumes that for the other incident types no tools need to be picked up.

5.3 Volunteer distribution

As the model contains thousands of volunteers, not every volunteer can be modeled independently. Instead, a distribution of volunteers is estimated, containing a certain probability ν_i that a random person is located in area i. Note that, as mentioned before, this probability is assumed to not be time-dependent. The expected number of available volunteers in area i can then be calculated as $\mu_i = \alpha n \nu_i$, where n is the total number of volunteers and α represents the fraction of volunteers that is available, i.e. they are present in the area and they are able to respond to a call. Also for α, an argument could be made that it requires time-dependency, as generally fewer people are available during nighttime. However, as this data is not really available and for most of the experiments it is not of importance, this is again assumed to be time-independent. When n is large enough, the distribution of volunteers can be modeled as a Poisson point process[11]. This is even the case when the individual volunteers have a unique availability probability and location distribution. As the number of volunteers is in the order of thousands, it is assumed that the properties of the Poisson point process can be used in further calculations regarding the location of volunteers.

For the spatial distribution ν, there are two natural options. The first is to make it equal to the spatial distribution of population. This is the case when every person has an equal probability of being a civil volunteer. Note that this is based only on the residence of civilians, so it is not entirely realistic when also taking into account movement to jobs, etc.

The second option is to optimize the distribution by minimizing the late responses. Late responses are defined as responses where the target time τ of 8 minutes. Although the FDAA does not have direct influence on this distribution and reaching this exact distribution is therefore not realistic, it does give a good indication of where to focus the efforts of finding volunteers. The optimization of this distribution is not a main focus of this research, as it had already been researched previously. However, an intuitive overview of how the optimization works will be given in the following section. For more technical details, refer to *van den Berg (2021)* [11] or *Schönberger (2021)*[1].

The first step of the process is to divide the region in multiple subregions. These subregions are the same neighborhoods that were used for the incident rates in Section 5.1. The distance that the volunteer can cover within this target time is defined as d_τ. $T_\nu(i)$ is a random variable representing the time it takes for a volunteer in subregion i to reach an incident under the volunteer distribution ν. Then, for each region, the probability that a volunteer will arrive within a target time is calculated:

$$\mathbb{P}(T_\nu > \tau) = \sum_{i \in I} \lambda_i \mathbb{P}(T_\nu(i) > \tau) = \sum_{i \in I} \lambda_i \exp(-\pi d_\tau^2 n \alpha \nu_i / a_i)$$

The last part of the expression can be inferred by using the total area of the circle with radius d_t to calculate the total density multiplying this by μ_i divided by the area of subregion i. This sum is then used as an objective function, essentially maximizing the probability that for any incident, the closest volunteer is at most a distance of d_τ from the incident. Note that this is not the same as directly minimizing the expected response time of the closest volunteer to any incident. This problem can then be solved by a greedy algorithm by adding a small amount of volunteer mass to the region with the highest marginal benefit. This marginal benefit is calculated by taking the absolute value of the derivative with respect to ν_i of the objective function for location i only:

$$\begin{aligned}
&\mathbb{P}(T_\nu(i) > \tau) = \lambda_i \exp(-\pi d_\tau^2 n \alpha \nu_i / a_i) \\
&|\frac{df}{d\nu_i}| = |-\lambda_i \pi d_\tau^2 n \alpha / a_i \exp(-\pi d_\tau^2 n \alpha \nu_i / a_i)| = \lambda_i \pi d_\tau^2 n \alpha / a_i \exp(-\pi d_\tau^2 n \alpha \nu_i / a_i)
\end{aligned}$$

5.3.1 Sampling from the volunteer distribution

In a certain sub-area B in a Poisson point environment, the number of volunteers in this sub-area follows a Poisson distribution with parameter $\lambda = \frac{\alpha n \nu_i}{a_i} * a(B)$, where a_i represents the area of region i and $a(B)$ represents the area of region B. Note that from this equation, the same formulation for μ_i can be inferred as in the previous section, by replacing B by region i and using the fact that the expectation of the Poisson distribution is equal to its parameter λ. In later parts of this thesis, a circle with radius r is used as region B, meaning $\lambda = \frac{\alpha n \nu_i}{a_i} * \pi r^2$. To then sample the distance of the volunteers, a technique called inverse transform sampling is used.

Suppose there exists a circle with radius r in which the density of volunteers is distributed uniformly. In the circle, we can draw a smaller circle with radius d ($0 \leq d \leq r$), see Figure 5. In this example, we take D as the random variable, as this distance is what we want to sample. If we sample a volunteer in this area, we can assume that the total probability mass in the big circle is 1. The Cumulative Distribution Function (CDF), which is by definition $\mathbb{P}(D \leq d)$, can be calculated by dividing the area of the circle with radius d by the total area: $\text{CDF}(d) = \frac{\pi d^2}{\pi r^2}$. For inverse transform sampling we need the inverse of the CDF, which yields:

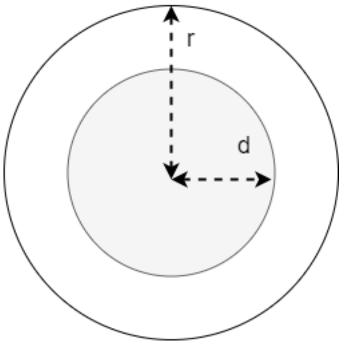

Figure 5: *Circles with uniform volunteer density*

$$\begin{aligned}
\mathrm{CDF}(d) &= \frac{\pi d^2}{\pi r^2} \\
x &= \frac{\pi d^2}{\pi r^2} && \text{represent CDF with variable } x\ (0 \leq x \leq 1) \\
d^2 &= x * r^2 \\
d &= r\sqrt{x} && \text{Now } x \text{ and } d \text{ have been swapped} \\
\mathrm{CDF}^{-1}(x) &= r\sqrt{x}
\end{aligned}$$

The distance from the incident given that the volunteer is in a circle with radius r around the incident can then easily be sampled by plugging a random uniform variable x $(0 \leq x \leq 1)$ in the formula $d = r\sqrt{x}$. Note that this technique of sampling assumes a uniform volunteer density in the circle around the incident. This assumption is generally valid, but for areas with very low density, the closes volunteer might in practice be in a different area with also a different volunteer density.

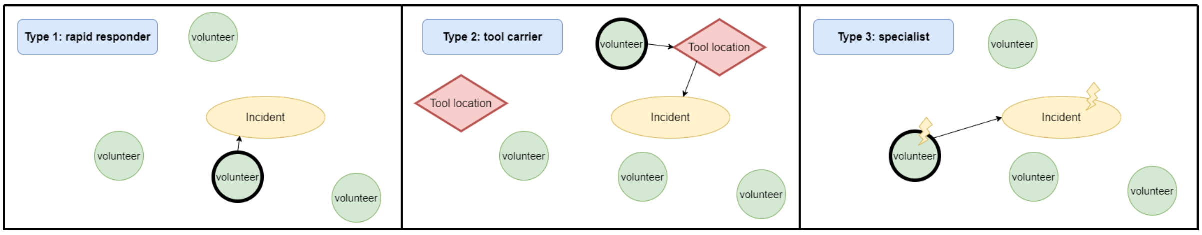

Figure 6: *Graphical representation of the different volunteer types*

5.4 Volunteer types

In the models that consider the civil volunteers, three different types of volunteers are distinguished, namely:

1. Rapid responders
2. Tool carriers
3. Specialists

Important to note is that the volunteer type is not predetermined for every volunteer: for some incidents, a person can be a rapid responder, whereas for a different incident they could be a tool carrier or a specialist. As such, the volunteer type is more related to the role of the person when being called for an incident. Figure 6 shows a graphical overview of the different volunteer types.

Rapid responders are volunteers that go directly to an incident, without making a detour to for example pick up tools. For any model that considers civil volunteers, rapid responders are always present. Because they are the first arriving volunteer, they are able to quickly gather preliminary information that the FDAA and other volunteers can use to more accurately assess the situation. Another example of their possible tasks is clearing the area around the incident of bystanders so the fire department can approach the incident more smoothly. Typically, people closest to the incident location are assigned the task of rapid responder, as they are able to arrive to the incident the fastest.

Tool carriers are volunteers that first pick up tools from a tool location before travelling to the incident. This means that they will arrive later than rapid responders, so their main added value is the tools they bring to the incident. Which tool this might be depends on the specific type of incident, but examples are AED's, security tape, or perhaps even fire extinguishers for small fires. As fire trucks often equipped with all the tools needed for an incident, tool carriers only have added value when they can arrive to an incident location before the FDAA. Tool carriers do not necessarily respond to each incident: one can choose to let them only respond to certain incident types.

Specialists are volunteers that have a certain kind of expertise in an area useful to the FDAA. Examples of this include electricians and chemical experts. Although fire fighters are trained to deal with such incidents themselves, this is only a basic training and therefore an expert in this subject could be very useful if the situation becomes too complex or dangerous. This is especially true when this type of incident is rare or very new, meaning that the fire fighter has not much practical experience with dealing with this type of problem. Currently, an expert from the fire department is called in these situations, but this is only one person in the whole province meaning that travel times are very long, especially when this one person is unavailable. In this regard, volunteers can already add a lot even when there are only a few people with a certain expertise in the population. When applying for the program, it would be useful if volunteers can register in which area they have a useful expertise, if any. As long as they can respond to an incident within an hour or so, they are already of added value to the FDAA. The model always includes travel times for specialists, but in practice video-calling can also be a viable option. This is an important detail when interpreting the results.

For ease of modeling, the volunteer types are assumed to be entirely distinct. This means that a volunteer can not have two tasks at once, e.g. being both a rapid responder and a specialist. In practice, this assumption is not very restricting since specialists are assumed to be only a very small part of the volunteer population. Additionally, when a volunteer is a tool carrier, it is not likely that they are also the first person to arrive to an incident as they will also need to pick up the tools.

5.5 Tool placement

Strategic placement of tools is one of the main subjects of this thesis and also one of the main contributions compared to other research on the topic of civil volunteering, and specifically compared to the previous research for the FDAA by *Schönberger (2021)*[1]. The locations of certain tools are placed optimally such that they meet as much demand as possible, while also restricting the total number of tool locations used. All models use the KRO (Kern Registratie Objecten) dataset[1] to select the public buildings that are suitable for placing tools. It is very likely that not all those locations are in fact suitable, but they do make a good impression on how many locations are available and where they are located over the safety region. These locations are plotted in Figure 7a. An interesting detail is that there are also locations outside the drawn region. These buildings are located in the area of Weesp, which was added to the municipality of Amsterdam and therefore to the safety region Amsterdam-Amstelland as of March 24 2022. However, as there is barely any data regarding the number of incidents in this area, it is not taken into account for this research, meaning that locations in this area will not show up in the results. Nevertheless, it is important to remember that in practice this area should be included. Another important detail is that many areas feature no potential location at all. From a practical standpoint this makes sense, because these areas are mostly lakes, port areas, or areas with very low population density and economic activity. This does mean that it is impossible to cover all potential incident locations, as many of them have no public buildings nearby.

[1]This dataset is one of the internal datasets from the FDAA. However, similar data can be found online, such as on the website of the municipality: https://maps.amsterdam.nl/open$_{g}$*eodata*/.

The mathematical model used to represent the tool placement problem is largely based on the model from *van den Berg et al. (2017)* [5], but with the context of placing tools instead of fire stations. In this model, demand is separated in different demand locations, with each location creating demand from its centroid. Distances from tool locations are then calculated based on this centroid. The main downside of this assumption is that in the optimal solution, large neighborhoods will have only 1 tool location close to the center. This means that the optimal solution can calculate fewer tool locations than are actually needed. However, since most of the regions, especially in the city center, are small enough, this is not a big problem.

5.5.1 Simple tool location problem

In the first problem, the placement of tools is optimized in the most basic way, without making a distinction between different types of incidents or tools. The objective function maximizes the number of locations covered, weighed by the demand of that location. In this case, a location is covered if the volunteer can travel from the tool location to the demand location (which is taken as the centroid of a neighborhood) within 8 minutes. The objective also attributes a cost β to opening new locations, so that the number of tool locations will not go to infinity. Constraint 1a assures that a location can only be covered if a tool is located at at least one of the locations that are reachable within the given maximum response time.

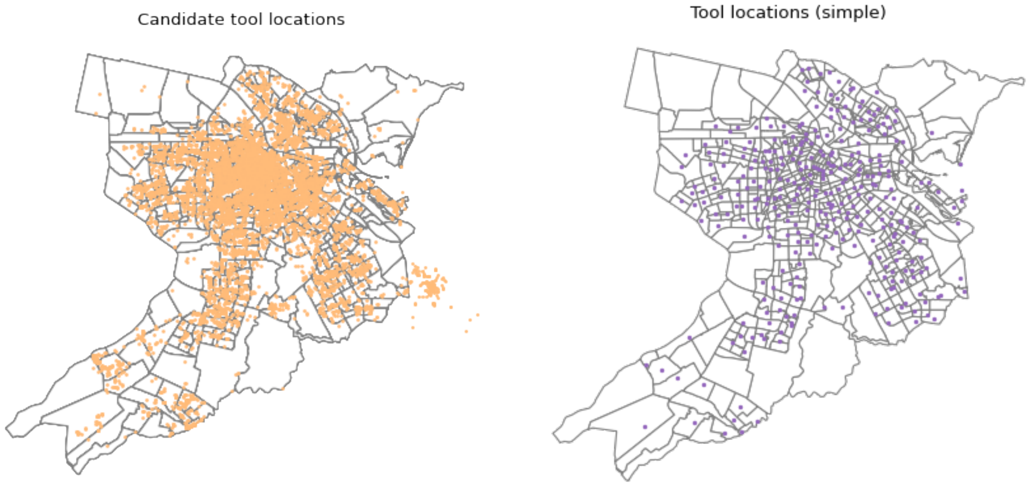

(a) *Map of candidate tool locations*

(b) *Result of basic tool optimization model*

Figure 7: *Plots of candidate and optimal tool locations*

Input

- N: Set of demand locations (neighborhoods)
- M: Set of potential tool locations (public buildings)
- d_i: Demand at demand location $i \in N$
- r_i: Response time target for demand location $i \in N$. Taken as 8 minutes for every demand location.
- t_{ij}: Travel time from location $j \in M$ to demand location $i \in N$. Calculated using the OSRM [14] distance and an average travel speed of 5 km/h.
- β: Penalty on number of opened base locations (in number of uncovered incidents in the data)

- M_i: Set of tool locations that can cover demand location $i \in N = \{j \in M : t_{ij} \leq r_i\}$

Variables

- y_i: Binary variable indicating whether demand point i is covered by a tool location
- z_j: Binary variable indicating whether a tool is located at tool location j

Model

$$\begin{aligned} \max \quad & \{\sum_{i \in N} d_i y_i - \beta \sum_{j \in M} z_j\} \\ \text{s.t.} \quad & \sum_{j \in M_i} z_j \geq y_i \quad \forall i \in N \qquad (1a) \\ & y_i, z_j \in \{0, 1\} \quad \forall i \in N, j \in M \end{aligned}$$

5.5.2 Multi-type tool location problem

In the multi-type setting, only incidents of the types in set K are taken into account. Each of these types has its own spatial distribution and also its own tools that can only be used for incidents of that specific type. This means that it is assumed that incidents of other types do not require tools and are therefore not considered in the response times. Variants where one tool can be used for multiple types are not considered, but require only minor changes to the model. Each tool has its own cost δ and each location has an opening cost of β, which is independent of the number of tools placed at that location. This way, tools are incentivized to be grouped together rather than opening a new location for every new tool type. The implementation of these costs is similar to the one of the simple tool placement and can be seen in the objective function of the model below. All covered demand is subtracted by the costs of all tool locations and the costs of all individual tools. Constraint 2a again makes sure that a location can only be covered when a nearby tool is present. Constraint 2b implements the behavior of variable z_j, where it must be 1 if a tool of any type is present at a location.

Input

- N: Set of demand locations (neighborhoods)
- M: Set of potential tool locations (public buildings)
- K: Set of incident types and corresponding tool types
- d_{ik}: Demand at demand location $i \in N$ from incidents of type $k \in K$
- r_i: Response time target for demand location $i \in N$. Taken as 8 minutes for every demand location.
- t_{ij}: Travel time from location $j \in M$ to demand location $i \in N$. Calculated using the OSRM distance and an average travel speed of 5 km/h.
- β: Penalty on number of opened base locations (in number of uncovered incidents in the data)
- δ: Penalty on the total number of tools used (regardless of type). Can be either a single value or an array of values where δ_k represents the penalty for tool type $k \in K$
- M_i: Set of tool locations that can cover demand location $i \in N = \{j \in M : t_{ij} \leq r_i\}$

Variables

- x_{jk}: Binary variable indicating whether there is a tool of type k located at tool location j

- y_{ik}: Binary variable indicating whether demand point i is covered by a tool of type k
- z_j: Binary variable indicating whether any tool is located at tool location j

Model

$$
\begin{aligned}
\max \quad & \{\sum_{i \in N} \sum_{k \in K} d_{ik} y_{ik} - \beta \sum_{j \in M} z_j - \sum_{j \in M} \sum_{k \in K} \delta \text{ (or } \delta_k) \; x_{jk}\} \\
\text{s.t.} \quad & \sum_{j \in M_i} x_{jk} \geq y_i k \quad \forall i \in N, k \in K && \text{(2a)} \\
& x_{jk} \leq z_j \quad \forall j \in M, k \in K && \text{(2b)} \\
& y_i, z_j \in \{0, 1\} \quad \forall i \in N, j \in M, k \in K
\end{aligned}
$$

5.6 Volunteer responses

For a volunteer to respond to an incident, they first need to be notified of the incident and the incident location. In this thesis, it is assumed that volunteers are notified via an app that is able to acquire the location of volunteers when an incident takes place nearby. Obviously this is only possible with the necessary privacy considerations. An alternative to this system is to use the home or work address of volunteers and notify them when an incident occurs near this address. This practice is safer from a privacy standpoint, but it severely reduces the effectiveness from notification meaning that more people will need to be alerted and the closest volunteer can not always be found. For this reason, civil volunteer programs are usually steering away from this method in favor of an app in which one's location is tracked [12].

Each notified volunteer has a probability α to accept the call and also a pre-trip delay of τ. This pre-trip delay includes all the time between receiving the notification and leaving the building. After this, the volunteer will move towards the incident with speed w. If the volunteer is a tool carrier, they will first travel to a tool location and then towards the incident location. Volunteers of other types will travel directly to the incident location taking the fastest route possible.

For the purpose of easy sampling, response times of volunteers are calculated using direct distances instead of actual distances. To compensate for this, the distance is divided by a factor of 1.3, as a factor between 1.2 and 1.4 is advised by the Ministry of Infrastructure and Environment (2015) [15].

5.6.1 Smart notification

Alerting the correct people is a challenge in itself, even when all their locations are known. Typically it will not be enough to only alert the closest volunteer, as volunteers are not always available and it is not known beforehand whether someone will be available or not. In practice, the k closest volunteers or all volunteers within an area of t meters around the incident location are alerted. Besides this, it is also important which type to attribute to which volunteer. Lastly, deciding on how many volunteers to alert is also an important task. This number should not be too high, as many people will then be called for nothing, but alerting a low number of volunteers can also be risky as no-one might show up. For these reasons, smart notification will be one of the main challenges tackled in this thesis.

The acceptance probability α is taken in literature to be 5 or 10 %. This number is indeed quite low, even when taking into account that people are often busy or sleeping. The expectation is that this number will increase by smart notification, as it will give people the idea that they are only notified when they are really needed. This can however not be taken for granted and for this reason, α will vary between 5 and 10 % in most scenarios that are considered in this research. The effect of smart notification on the acceptance probability is therefore something that can only be tracked once the civil volunteer program is actually deployed. Besides acceptance probability, there exists a concept called *acceptance delay*, meaning that not all volunteers will accept or deny the task immediately. This acceptance delay is assumed to be constant and is taken as a part of the pre-trip delay τ.

The first and main method of alerting that will be used in this research is called *k-closest alerting*. In this method of alerting, the k closest volunteers are alerted. Here, an important distinction is made between volunteers of type 1&3, and volunteers of type 2. For rapid responders and specialists, the k closest volunteers to the incident location are called. For tool carriers, the k volunteers that have the shortest distance to the incident via a tool location are called. This is calculated by finding all reasonably close tool locations and sampling k volunteers around every tool locations. Of all those volunteers, the k volunteers with the shortest total distance are selected.

With this method, it is possible that no volunteers at all are available. If this is the case, the next k closest volunteers will be notified. As there is a significant time difference between the first round of notifications and the next round of notifications, a time of ϵ will be added to the pre-trip delay. ϵ can be seen as roughly the acceptance delay, since the extra time between the two notification rounds consists of waiting for all the volunteers to accept or dismiss the notification.

The second method is called *area alerting.* In this method, all volunteers inside a radius will be called. Again, this calculated either from the incident location or candidate tool locations, depending on the volunteer type. When no available volunteers are present within the area, the radius of the circle is extended by 50%, meaning that the total area will have size $\pi * (1.5r)^2 = 2.25\pi r^2$, instead of πr^2. Obviously, only the volunteers in the extra area will be alerted, since the volunteers in the original area all dismissed the notification. Since there is no strict bound on the number of volunteers alerted per round, the same technique is used for all candidate tool locations. The more technical implementation of sampling from the distribution is explained in Section 5.3.1.

6 Experiments

In this section, the experimental setup of this research is described. These experiments take place in the model that was described in Section 5. For all experiments, the reasoning behind several choices is explained as well as how it relates to the research questions that were stated in Section 2. Table 1 shows an overview of how all parameters are changed between the different experiments. In Section 6.9, a description of all the parameters is provided.

Most of the scenarios differ by (a combination of):

- Changing important parameters such as the number of volunteers or the volunteer speed
- Changing complexity of the model, such as making a distinction between incident types or adding an extra volunteer type
- Changing the notification policy by restricting or changing the number of people called, or by making other smart policy choices

Where applicable, the model was trained or optimized using historical data. The calculations of the results were done based on data generated by the simulator, as explained in Section 4.4.

6.1 Experiment 1: simple model

In the first experiment, response times are calculated only for the simplest scenario, meaning the volunteers can only be rapid responders or tool carriers and no distinction is made between the incident types. For this reason, tool locations are determined by solving the simple tool location problem, meaning no distinction is made between different types of tools. The volunteers are called via area notification with a starting radius of 100 meters. This radius was chosen so that it is most comparable with the other strategy in terms of number of people that receive a notification. Other scenarios are derived from this simple model by changing one or more parameters.

6.2 Experiment 2: optimized volunteer distribution

The second experiment focuses on the optimized volunteer distribution, as is discussed in Section 5.3. Although this experiment is virtually equivalent to the research done in *Schönberger (2021)* [1], the results could be significantly different as some important calculations were changed[2]. Moreover, the implementation of the notification system was changed to simulate more realistically the time to find the closest available volunteer. For ease of comparison, all other hyperparameters were set the same as in experiment 1.

6.3 Experiment 3: k-closest notification

Experiment 3 tests the alternative alerting method of k-closest notification. This also relates to research question 2: *"How can volunteers be notified in a smart way so that they are only sent out when they are really needed?"* and research question 3: *"What are the benefits and costs of smart notification?"*

6.4 Experiment 4: multi-type tool placement

The main goal of experiment 4 is to make a distinction between incident types and to provide and answer to research question 5: *"How can we distinguish between different types of tools and how will this impact the response time of tool carriers?"* Although from a modeling perspective it is not very important, some types were chosen that are likely to require tools. These types are: {Fire, Emergency service, Healthcare, Environment, Accident}. To have a better view on the impact of the tool costs, they were set at very

[2] When implementing the optimized volunteer distribution into the new model, some significant bugs were discovered when calculating the optimal volunteer densities. The impact of these bugs is difficult to estimate.

Parameters	n	α	Vol. types	Strategy	Radius	k	Tool costs	Distribution
Experiment 1	100,000	0.1	1,2	Area	100 m	-	-	Proportional
Experiment 2	100,000	0.1	1,2	Area	100 m	-	-	Optimal
Experiment 3	100,000	0.1	1,2	k-closest	-	20	-	Proportional
Experiment 4	100,000	0.1	1,2	Area	100 m	-	100, 50, 20, 1, 10	Proportional
Experiment 5	100,000	0.1	1,2,3	k-closest	-	20	-	Proportional
Experiment 6	50,000	0.1	1,2	Area	100 m	-	-	Proportional
Experiment 7	100,000	0.2	1,2	k-closest	-	20	-	Optimal
Experiment 8	100,000	0.2	1,2,3	k-closest	-	20	100, 50, 20, 1, 10	Optimal

Table 1: *Overview of parameters used in different experiments*

different but still realistic values. All the exact parameter values can again be found in Table 1. It is quite difficult to compare the results of this experiment with the results of experiments with simple tool placement: firstly because the results are presented per type and are difficult to interpret as a whole, and secondly because the number of tool locations is not comparable since not all tool locations have all tool types present.

6.5 Experiment 5: specialists

In experiment 5, the volunteer type of expert was added. The number of specialists was set to 1% of the total population of volunteers, which in this experiment results in 1000 specialists. Although in practice specialists would only respond to very specific types of incidents, the specialists in this experiment respond to all incidents. The reasoning behind this is that this distinction does not have any real influence on the results, but it does give much more data which makes it possible to better evaluate the performance of the specialists. As the density of specialists is relatively low, k-closest was used as notification strategy, because otherwise many call rounds would be needed before finding an available specialist.

6.6 Experiment 6: 50,000 civil volunteers

In this experiment, the number of volunteers in the network is halved. This should give an indication on how many volunteers are needed and what the consequences are when the goal of 100,000 volunteers is not reached. Besides this, none of the basic settings from experiment 1 are altered.

6.7 Experiment 7: best-case scenario

In experiment 7, a best-case scenario within the model assumptions is created by optimizing the volunteer distribution, as described in Section 5.3, and also using the k-closest notification, as this is likely the more efficient way of notifying volunteers compared to area notification. In this context, 'efficient' means that the policy does not necessarily minimize all the response times, but rather makes a good trade-off between keeping response times low and not sending too many notifications. As an effect of smarter notification strategies, the acceptance rate is doubled in this scenario. The reasoning behind this is that people are likely more willing to accept a call if they know that they will indeed be needed at the incident, instead of 5 other people showing up too. Even though there is no dedicated worst-case scenario in the experiments, the scenarios with basic settings such as area notification do serve this purpose.

6.8 Experiment 8: best-case scenario, multi-type setting

Experiment 8 adds the dimension of multi-type tool placement to experiment 7. Equivalently, it adds optimization of the volunteer distribution and k-closest notification to experiment 3. Furthermore, specialists are also added to the model to see how their performance is influenced by better policies.

Parameter	Explanation	Value
Volunteer speed	Speed with which the volunteer travels in the network	5 km/h
Target time	Used for volunteer distribution and tool optimization	8 min
β	Opening cost of new demand location (in number of incidents covered)	10
τ	Pre-trip delay	2 min
ϵ	Notification delay	1 min
Specialist %	Percentage of volunteer population that is a specialist (if applicable)	1%

Table 2: *Overview of hyperparameters that are unchanged over the experiments*

6.9 Parameters used

As mentioned previously, Table 1 shows the parameters that are important for the distinction between experiments. These parameters are, in order:

- The number of volunteers in the program
- The acceptance rate α
- The volunteer types used (1 for rapid responder, 2 for tool carrier, and 3 for specialist)
- The deployed notification strategy
- The radius of the circle in which notifications are sent out in the first call in the area notification strategy
- The value of k used in the k-closest notification strategy
- The percentage of the volunteer population that is a specialist
- The costs of tools with the respective type. The exact tool types are (in this order): {Fire, Healthcare, Accident, Emergency service, Environment}.
- Calculation of the volunteer distribution (optimal or proportional)

Besides these parameters, there are some parameters that have the same value for each experiment. This was done either because there is not enough practical value to experiment with the parameter values, or because the influence of this parameter was already tested in *Schönberger (2021)* [1] and it was deemed unnecessary to test it again. For completeness and academic transparency, these parameters and values are listed in Table 2. Note that, although the choice of parameters is important for this research, similar experiments can be performed with different parameter values if those values turn out to be more realistic.

7 Results

In this section, the results of the experiments as described in Section 6 are presented. First, the different evaluation metrics are discussed, as well as the reasoning behind the choices for these metrics. Next, the results of all the individual experiments are shown, as well as a comparison across multiple experiments. Lastly, an analysis of the run times is performed in Section 7.5.

7.1 Evaluation metrics

The main way in which the results are presented is in a colored map, where the colors stand for the mean difference between the arrival time of the volunteer and the arrival time of the first fire truck, like in Figure 9. This so-called "time gain" is the most important metric for rapid responders and tool carriers. The performance of specialists is instead reported by their response time, as they also have value when arriving to the incident after the fire department.

Besides this graphical evaluation, differences between two related experiments are also tested statistically. This is either done via the paired t-test, or via the Wilcoxon signed rank test if the differences do not follow a normal distribution. Normality is tested with the Shapiro-Wilk test. Due to the high sample size, statistical tests can detect even small differences from normality when this is not quite necessary. For this reason, the differences between the two samples are also plotted with a histogram to give a visual confirmation of the conclusion of the test. For rapid responders and tool carriers, the difference in their mean time gain is tested, and for specialists the difference in mean response time is tested.

Thirdly, a table is provided with all the relevant statistics for all experiments. This includes mean and median time gain/response time, fractions of incidents responded to faster than the fire department, and the average number of people that received a notification. These metrics should give a good view of the performance of a policy from different sides, as well as the trade-off costs of the number of people notified per incident.

7.2 Tool placement

In Figure 7b, the results from the basic tool optimization model were already shown. Figure 8 shows the chosen tool locations of types "Brand" (fire) and "Dienstverlening" (emergency service). These two types had respective tool costs of 100 and 1 units. This measurement can be roughly translated as the number of incidents of that type that one tool location covers. They are also the two most common incident types as shown in Section 4. For these reasons, these types were selected for comparison throughout this section. As is visible in the picture, many tool locations are shared between the two types, which makes sense as this reduces costs. On first sight, there does not seem to be a large difference between the number of locations of the two types, but after some inspection it becomes apparent that there are much less tools for type Fire. This also makes sense, as those tools have a much higher cost. If anything, the effect might actually be less than expected. Also the fact that many parts of the region have no candidate locations is clearly visible in the result. Some parts on the edge of the region have no or barely any tool locations. Tool locations of the other types can be found in the appendix.

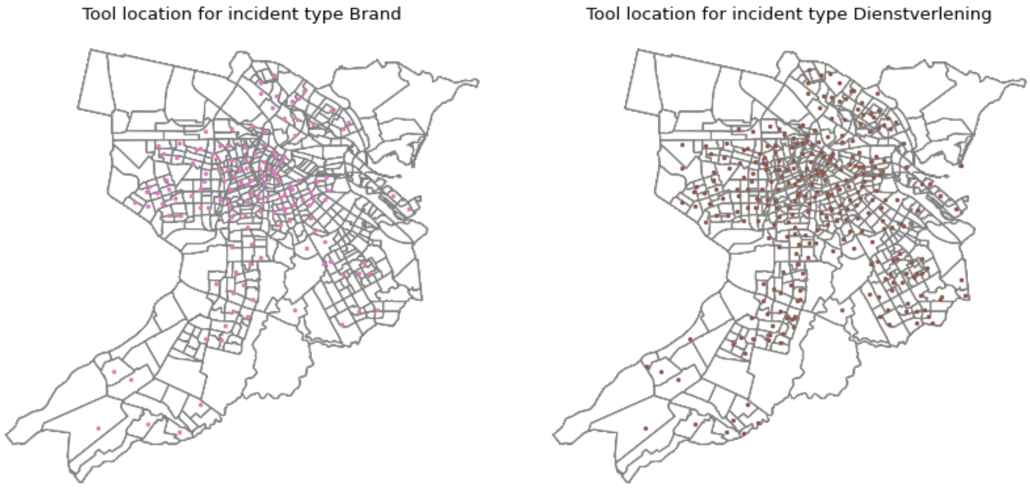

(a) *Incident type "Brand"* **(b)** *Incident type "Dienstverlening"*

Figure 8: *Maps of tool locations for different incident types*

7.3 Experiment results

In this section, the results of all individual experiments are presented. This includes the colored maps and a table with the most important statistics at the end of the section. Additionally, some small remarks about the interpretation of the results and the comparison between experiments will be given.

7.3.1 Experiment 1: simple model

Figure 9 shows the results from experiment 1, which is the starting situation for all experiments. Immediately it becomes apparent that the time gains are much greater for rapid responders. This is by definition the case, as they can respond to the incident immediately instead of first having to walk to a tool location. The tool carrier response times do not seem to be too heavily affected by the lack of tool locations in some neighborhoods. In general, the response time gain seems to be around 2 to 6 minutes. In areas such as the Western port area, tool carriers surprisingly are faster on average than rapid responders. This effect is difficult to explain, but one reason for this might be that tool carriers are being called in a radius around all close tool locations, whereas rapid responders only are called around the incident location. For some neighborhoods, this increased number of people called might make the difference in response time.

A couple of areas show almost no time gain, meaning that the volunteer is almost never earlier present than the fire department. The most likely reason for this is that the area is very sparsely populated and therefore there are no volunteers close by.

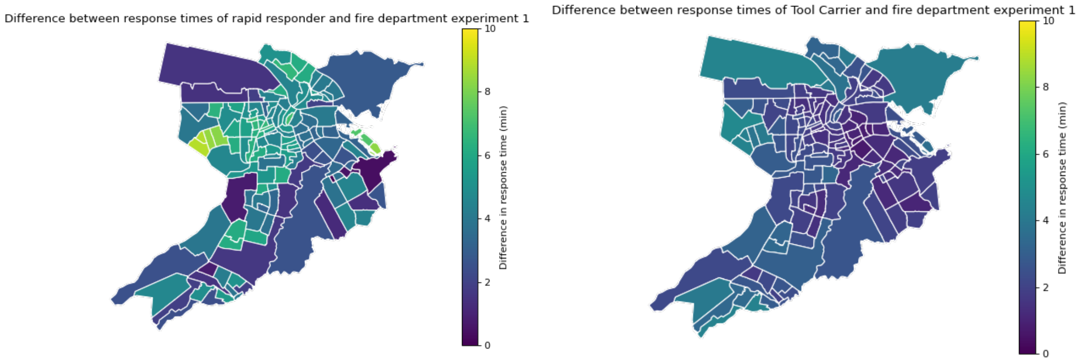

(a) *Rapid responder*

(b) *Tool carrier*

Figure 9: *Mean time gains for the different volunteer types of experiment 1*

7.3.2 Experiment 2: optimized volunteer distribution

In Figure 10, the results from experiments 2 are shown. When comparing this to experiment 1, some significant time gains are visible for the rapid responder in the lesser populated areas. For the central area, the time gains seem similar or perhaps slightly decreased. Apparently, the optimized volunteer distribution favors less populated areas compared to the proportional distribution. In general, the distribution of time gains seems to be much smoother, meaning there is more fairness across the region in terms of response time. For tool carriers, this difference is not very noticeable.

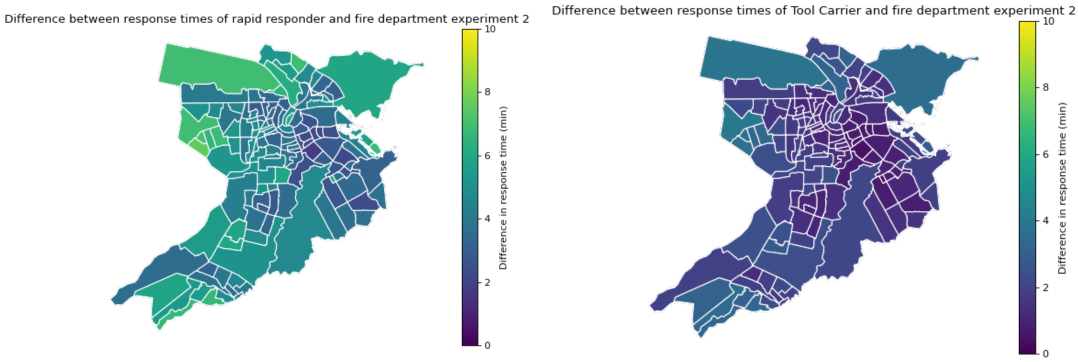

(a) *Rapid responder*

(b) *Tool carrier*

Figure 10: *Mean time gains for the different volunteer types of experiment 2*

7.3.3 Experiment 3: k-closest notification

Figure 11 shows the results from experiment 3, where k-closest notification was implemented. When comparing this to the results from experiment 1, some significant improvements have been made for both rapid responder and tool carriers. However, there is still quite a big difference in time gains between neighborhoods, which might be a problem depending on how many incidents it affects.

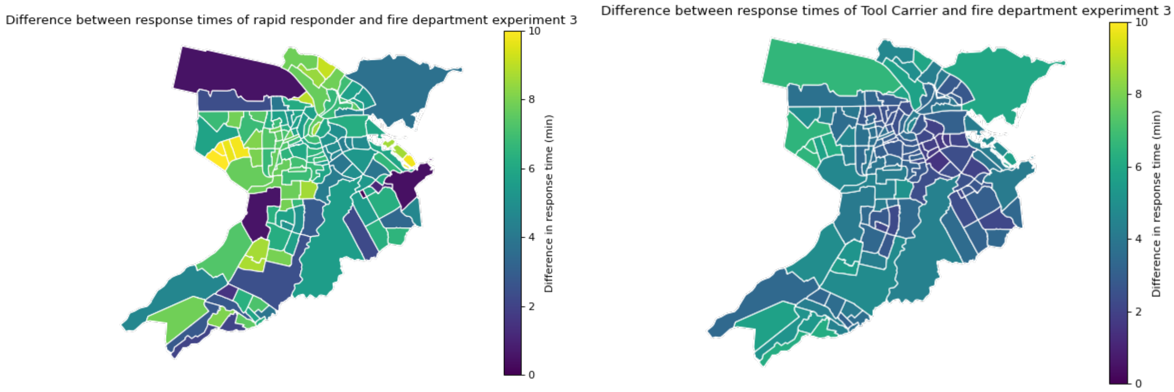

(a) *Rapid responder* **(b)** *Tool carrier*

Figure 11: *Mean time gains for the different volunteer types of experiment 3*

7.3.4 Experiment 4: multi-type tool placement

Figure 12 shows the time gains for the tool carrier-volunteers for the incident types Fire and Emergency Service. The difference between the two types does not seem to be large, but over-all incidents of type emergency service seem to have a lower time gain. This makes sense, as there are also more tool locations for this incident type. The results of the rapid responders are not shown, as they are the same as in experiment 1. The results from the other tool types can be found in the appendix

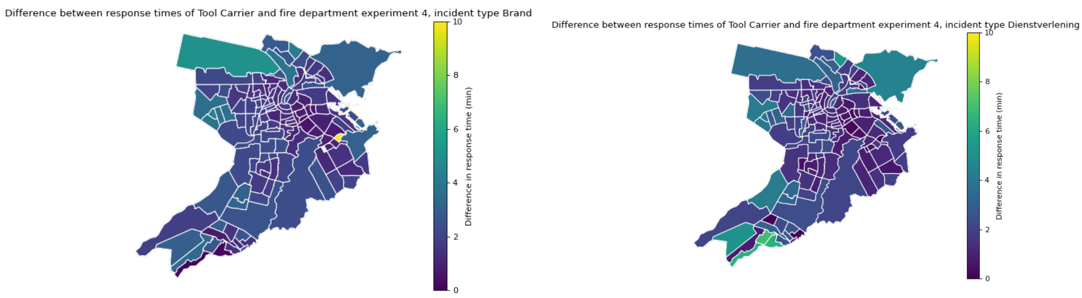

(a) *Fire* **(b)** *Emergency Service*

Figure 12: *Mean time gains for the different incident types of experiment 4 tool placement*

7.3.5 Experiment 5: specialists

Figure 13 shows the response times of the specialists in experiment 5. Very apparent is the unrealistically high response time in the Western port area. This is likely due to the low population in that area. This also shows the limits of the model assumptions, because when the density of an area is this low, it is not realistic to assume that the nearest volunteer is actually in that area.

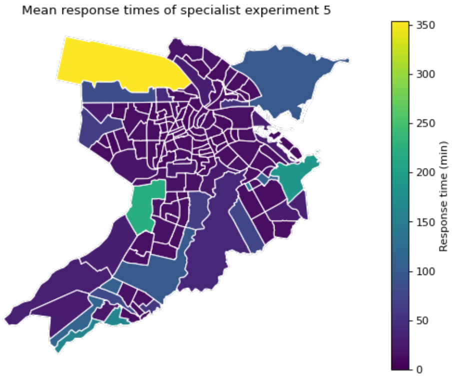

Figure 13: *Mean response time of specialist in experiment 5*

7.3.6 Experiment 6: 50,000 civil volunteers

Figure 14 shows the time gains when the total number of volunteers is halved. Obviously, this means a loss in response time compared to experiment 1, but this loss is not very clear for either volunteer type.

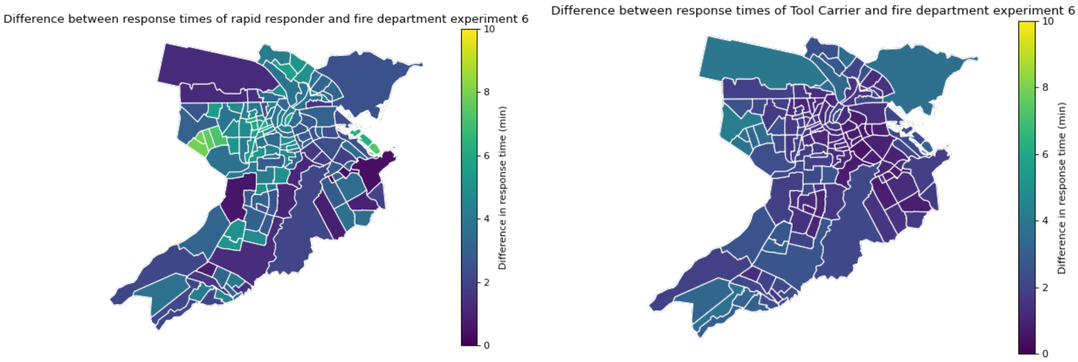

(a) *Rapid responder*

(b) *Tool carrier*

Figure 14: *Mean time gains for the different volunteer types of experiment 6*

7.3.7 Experiment 7: best-case scenario

In Figure 15, the results from experiment 7 are shown. The time gains on the map from this experiment are very large exceed those of previous experiments, especially for the rapid responder. It seems that for tool carriers the initial barrier of picking up tools makes it so that such improvements are difficult to achieve.

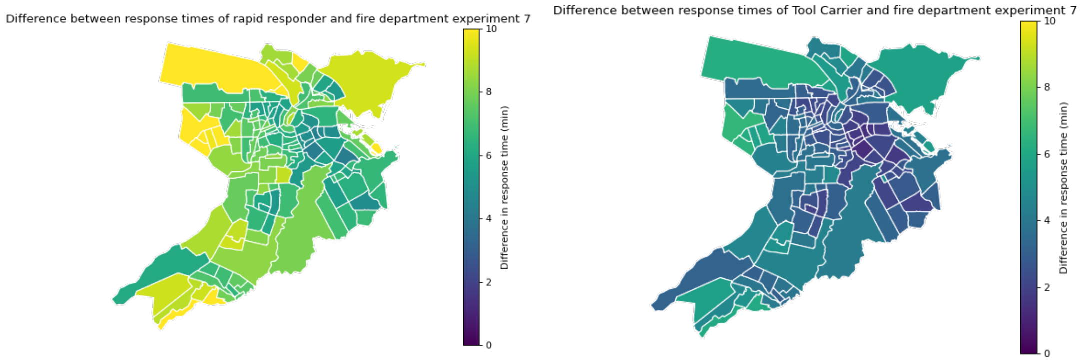

(a) *Rapid responder* **(b)** *Tool carrier*

Figure 15: *Mean time gains for the different volunteer types of experiment 7*

7.3.8 Experiment 8: best-case scenario, multi-type setting

Figure 16 shows the results from experiment 8, which was similar to experiment 7 but in a multi-type setting and with the added feature of specialists. Just as in experiment 4, there is not a clear difference in response times between the different incident types. The response time of the specialist seems to be better than that in experiment 5 (note the difference in axis), and mainly is much more evenly distributed. This is most likely due to the optimized volunteer distribution. A response time of around 3 minutes is very fast for a specialist and in many cases even faster than the fire department. The full results from the tool carriers are again visible in the appendix.

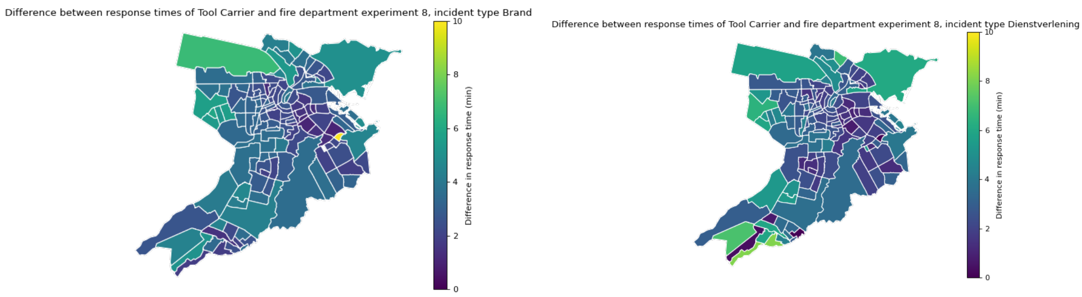

(a) *Tool carrier, incident type Fire* **(b)** *Tool carrier, incident type Emergency Service*

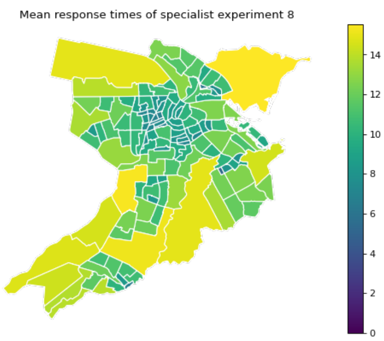

(c) *Specialist*

Figure 16: *Mean time gains for the different volunteer types of experiment 8*

7.3.9 General statistics

In this section, the statistics of the experiment results are presented and discussed. Tables 3 and 4 show the values for the rapid responders and tool carriers of the following statistics:

- **On time %**. The percentage of incidents where the volunteer arrived before the fire department
- **In target %**. The percentage of incidents where the volunteer arrived before the used target of 8 minutes
- **Mean time gain**, which was also presented regionally in the above sections
- **Median time gain**. This was chosen to make sure the results were not influenced too much by outliers.
- **Mean people called**. The mean number of people that received a notification to respond to an incident.

As can be seen in the tables, experiment 1 actually outperforms experiment 2 in terms of mean and especially median gain. This means that optimizing the volunteer distribution is not necessarily beneficial in terms of response time. A reason for this could be that most incidents are already near areas where people live and therefore the proportional distribution is already quite sensible. One important thing to note however is that optimizing the volunteer distribution does make sure that the network retains a good coverage over the whole region, which might be worth the loss in response time. A second reason is that experiment 1 had on average almost double the number of rapid responders and more than triple the number of tool carriers called. This is explained by the fact that there are more incidents in areas where the proportional distribution has a high volunteer density (i.e. it is a highly populated area). Obviously calling more people will increase the chances of finding someone that can respond immediately, which in turn decreases the response time.

Besides this conclusion, it is very apparent that k-closest notification is highly effective in both obtaining a good response time and also keeping the number of people called low. This can be seen by comparing the results of experiment 1 and 3 (proportional distribution), or by comparing experiment 2 and 7 (optimized distribution). Note however that experiment 7 also uses an increased acceptance rate, so this difference is not only caused by the notification policy. In general, the benefit from using the k-closest notification seems to be at least 1 minute in time gain, while also using much fewer volunteers. This is the case for both rapid responders and tool carriers.

The results from experiment 6 show that halving the number of volunteer does hurt performance significantly, but still the coverage remains quite good. The loss in performance is smaller than the gain in performance when moving from area notification to k-closest notification.

Lastly, the performance of rapid responders is very good, almost always arriving before the fire department and often with large time gains. This is not the case for tool carriers, although with k-closest notification their performance is also significantly improved. Specialists also perform quite well with an expected response time of approximately 10 to 16 minutes. Note that specialists typically do not need to arrive before the fire department, so comparison with other volunteer types is not logical. See Table 5 for the full results.

Statistic	RR on time %	TC on time %	RR in target %	TC in target %	RR mean gain
Experiment 1	87.5%	47.8%	89.3%	34.5%	286 s
Experiment 2	79.5%	38.7%	82.6%	19.5%	240 s
Experiment 3	96.6%	70.3%	96.9%	68.9%	382 s
Experiment 4	87.6%	47.5%	89.4%	34.1%	286 s
Experiment 5	96.6%	70.1%	96.8%	68,9%	382 s
Experiment 6	78.3%	40.3%	78.9%	23.7%	237 s
Experiment 7	99.7%	67.6%	100%	65.8%	415 s
Experiment 8	99.7%	66.4%	100%	63.9%	415 s

Table 3: *Relevant statistics of all experiments (1 of 3)*

Statistic	TC mean gain	RR median gain	TC median gain	RR mean called	TC mean called
Experiment 1	120 s	239 s	0 s	34	70
Experiment 2	93 s	182 s	0 s	19	19
Experiment 3	204 s	338 s	134 s	23	23
Experiment 4	119 s	238 s	0 s	34	71
Experiment 5	204 s	338 s	133 s	23	23
Experiment 6	97 s	179 s	0 s	22	45
Experiment 7	190 s	367 s	113 s	20	20
Experiment 8	186 s	366 s	107 s	20	20

Table 4: *Relevant statistics of all experiments (2 of 3)*

Statistic	Spec. in target %	Mean response spec.	Median response spec.	Mean people called spec.
Experiment 1	-	-	-	-
Experiment 2	-	-	-	-
Experiment 3	-	-	-	-
Experiment 4	-	-	-	-
Experiment 5	29.9%	1098 s	649 s	23
Experiment 6	-	-	-	-
Experiment 7	-	-	-	-
Experiment 8	36.6%	619 s	570 s	20

Table 5: *Relevant statistics of all experiments (3 of 3)*

7.4 Differences between experiments

In this section, the statistical tests to evaluate the difference in mean time gain (or response time) between two experiments are discussed. This leads to the two statistical hypotheses:

$$H_0 : \text{Sample x has the same expected value as sample y}$$
$$H_1 : \text{Sample x has a smaller expected value than sample y}$$

In all cases, the samples with a smaller sample mean were taken as sample x. The tested pairs of experiments are:

- Experiment 1 and 2, to verify the effect of optimizing the volunteer distribution.
- Experiment 1 and 3, to verify the effect of using k-closest notification instead of area notification. This was tested for both rapid responders and tool carriers, because smart notification is one of the main subjects of this research.
- Experiment 2 and 7, again to verify k-closest notification. For the same reason, both volunteer types were tested.
- Experiment 4 and 8, to test the performance of multi-type tool placement to what it can be in an optimal situation. Only tool carriers were considered.
- Experiment 5 and 8, to test the performance of specialists to what it can be in an optimal situation.

In all cases (except the first pair), experiment 1 and 2, the starting situation was improved on. All differences were tested to be non-normal and all pairs of sample had a statistically significant difference in mean, meaning all null-hypotheses were rejected (with significance level 0.05). An overview of the test can be found in Table 6. All further details of the results from the normality testing and difference testing can be found in the appendix. An interesting observation is that all the p-values from the tests were so low that they were essentially 0. This is likely due to the high sample size, giving the tests a high power.

Experiment pair	Volunteer type	Conclusion
1 and 2	Rapid Responder	The mean of experiment 2 is significantly smaller than that of experiment 1
1 and 3	Rapid Responder	The mean of experiment 1 is significantly smaller than that of experiment 3
1 and 3	Tool Carrier	The mean of experiment 1 is significantly smaller than that of experiment 3
2 and 7	Rapid Responder	The mean of experiment 2 is significantly smaller than that of experiment 7
2 and 7	Tool Carrier	The mean of experiment 2 is significantly smaller than that of experiment 7
4 and 8	Tool Carrier	The mean of experiment 4 is significantly smaller than that of experiment 8
5 and 8	Specialist	The mean of experiment 8 is significantly smaller than that of experiment 5

Table 6: *Results of pairwise testing*

7.5 Run time analysis

Besides the numerical results, the running times of the experiments and the initialization phases were also documented. Besides this, an estimation of the order on the number of operations was given, based on the number of input incidents (n). These measures are important to show how often and easily the calculations can be repeated, and how scalable the project is to include for example data of the whole country, instead of just one safety region. All results are shown in Table 7. Note that no single experiment requires all actions to be performed. For example, experiment 3 only needs to calculate the proportional distribution, simple tool locations, and the response times for rapid responders and tool carriers with k-closest notification. The initialization steps only need to be repeated once the used data is deemed to be no longer representative for the current situation.

As the table shows, normal experiments only take one to three minutes approximately, without initialization. These operations also scale linearly with the number of incidents, meaning they are very scalable for larger data sets. All calculations were performed on a virtual server with Intel Xeon Gold 6126 2.6GHz CPU.

Action	Run time	Order of operations (estimate)
Fitting predictor to incidents	44 min, 34 s	$O(n^2)$
Forecasting incident rates for all types	7 min, 33 s	$O(n)$
Initializing & training simulator	2 h, 53 min, 28 s	$O(n^2)$
Simulating incidents	1 min, 30 s	$O(n)$
Calculating proportional volunteer distribution	3 s	$O(1)$
Calculating optimal volunteer distribution	51 s	$O(n)$
Calculating simple optimal tool locations	6 s	$O(n)$
Calculating multi-type optimal tool locations	1 min, 38 s	$O(n)$
Calculating rapid responder times (area notification)	9 s	$O(n)$
Calculating rapid responder times (k-closest notification)	4 s	$O(n)$
Calculating tool carrier response times (simple, area)	35 s	$O(n)$
Calculating tool carrier response times (simple, k-closest)	54 s	$O(n)$
Calculating tool carrier response times (multi, area)	48 s	$O(n)$
Calculating tool carrier response times (multi, k-closest)	48 s	$O(n)$
Calculating specialist response times (k-closest)	5 s	$O(n)$
Calculating test statistics	1 s	$O(n)$

Table 7: *Run time estimations of different actions*

8 Conclusion

One of the main goals of research of this thesis was to investigate and evaluate new methods of notification. In existing research, only area notification was used or no form of notification was considered (i.e. the closest available volunteer was selected without any extra considerations). k-closest notification was found to be a suitable alternative to area notification, beating area notification in both time gain and number people called. The likely reason for this is that k-closest notification is much more flexible and adapts better to the density of volunteers in the area. For this reason, a form of k-closest notification is recommended to be used in the future civil volunteer program.

A second important subject is the addition of the volunteer types of tool carrier and specialist. Under the tested parameters and assumptions, the tool carrier had an average time gain of around 3 minutes compared to the fire department. Specialists could arrive to the incident after around 10 to 16 minutes. The aim of this report is not to advise whether these volunteer should be deployed, but rather to give an indication to the fire department what their added value is given certain assumptions. However, considering specialists can also be called up via a video connection, it seems that their value is always significant regardless of their response time. This added value is less obvious for tool carriers and will only become clear once there is certainty about which tools exactly are needed for which incident types.

Lastly, it is definitely clear that a volunteer network, no matter in which form, can be beneficial to the FDAA. Their inherent nature of fast responses as well as other benefits such as getting in contact with experts and spreading messages in society is a clear proof of this. This is even the case when smaller numbers of volunteers are reached, instead of the desired 100,000. In general, more gains can be made by smartly implementing the network than by recruiting more volunteers. This also corresponds with the conclusions from Schönberger (2021) [1].

Once the civil volunteering network will be implemented, it is important to make room for the aforementioned suggested additions to the system. One way to do this is by documenting all the relevant skills of volunteers that sign up to the program, so that those skills can be put into practice once the network is fully functioning. If an app is created for sending the notifications, it is important that the location of volunteers can be tracked in a safe but accurate way and that options can be added for picking up tools or calling for a certain expertise.

9 Discussion

In this section, some of the limitations of the research are discussed, as well as future research on this topic. Section 9.1 covers the possible limitations related to the model assumptions and Section 9.2 discusses possible opportunities for future research in the field of civil volunteering.

9.1 Limitations

Many of the assumptions from this research are discussed in Section 5. These assumptions make the model more simple and insightful, but also steer the model away from reality making the results slightly less realistic. This is an inevitable result from this method of research, but a limitation nevertheless. This is especially visible in situations where the assumptions reach the limit of validity, for example in the case of the expert. In areas where there was a low volunteer density, the expert response times exploded because it was assumed that the nearest volunteer was close enough to the incident location that its volunteer density is the same as that on the incident location. In this extreme case, this assumption did not hold. Still, there is no indication that any of the assumptions majorly affected the results and therefore the conclusions of this research.

Another possible pitfall is the use of simulated data. Although this simulator is trained on real incident data and in this way resembles this data, it is possible that there are systematic differences between the two data sources which can influence the results.

9.2 Future research

For further research, it is always possible to expand the model with more options of research. This could be a new method of notification, or perhaps a completely different policy, potentially driven by issues that come up when deploying the network in practice. The possibilities for new notification strategies are virtually endless and this thesis only scratches the surface in that regard. Some ideas could be to go into the notification process in more detail, by adding options to cancel notifications and adding randomness to the time in which one responds to a notification. Another direction might be to give preference to certain volunteers based not only on distance, but also other factors such as skills, their historical acceptance rate, or the number of incidents the volunteer has already responded to. This could also be combined with giving a single volunteer multiple roles, by letting them be both a rapid responder and a tool carrier, for example. All of these changes have also impact on the number of people that receive a notification, which likely has some effect on the acceptance rate. Once there is more information about this relationship, this could also be added to the model.

Furthermore, it could be possible to investigate how response times of specifically tool carriers can be improved. Although their gains are already quite significant, it could prove useful to expand the candidate tool locations to non-public places. Secondly, a dual optimization problem could be created to optimize the volunteer distribution together with the tool placement optimization, instead of separately.

Lastly, one might look into shifting the main focus of the volunteer network from quickly responding to all incidents to assisting the fire department in areas where response times are currently long. This would create an extra net of safety and more evenly distribute the fire department coverage over the whole safety region.

References

[1] D. Schönberger, "Reducing the response time of fire department amsterdam-amstelland," M.S. thesis, Vrije Universiteit Amsterdam, 2021.

[2] J. M. Chaiken and J. E. Rolph, *Predicting the Demand for Fire Service*. Santa Monica, CA, 1971.

[3] G. Legemaate, J. de Deijn, S. Bhulai, and R. van der Mei, "Severe weather-based fire department incident forecasting," 2021.

[4] J. van den Bogaert, "The data-driven fire department," 2019.

[5] P. L. Van den Berg, G. A. Legemaate, and R. D. Van Der Mei, "Increasing the responsiveness of firefighter services by relocating base stations in amsterdam," *Interfaces*, vol. 47, no. 4, pp. 352–361, 2017.

[6] D. Usanov, G. Guido Legemaate, P. M. van de Ven, and R. D. van der Mei, "Fire truck relocation during major incidents," *Naval Research Logistics (NRL)*, vol. 66, no. 2, pp. 105–122, 2019.

[7] "Goodsam." (Apr. 2022), [Online]. Available: `https://www.goodsamapp.org/`.

[8] "Pulsepoint." (Apr. 2022), [Online]. Available: `https://www.pulsepoint.org/`.

[9] T. C. Chan, D. Demirtas, and R. H. Kwon, "Optimizing the deployment of public access defibrillators," *Management science*, vol. 62, no. 12, pp. 3617–3635, 2016.

[10] C. L. Sun, D. Demirtas, S. C. Brooks, L. J. Morrison, and T. C. Chan, "Overcoming spatial and temporal barriers to public access defibrillators via optimization," *Journal of the American College of Cardiology*, vol. 68, no. 8, pp. 836–845, 2016.

[11] P. L. van den Berg, S. G. Henderson, C. Jagtenberg, and H. Li, "Modeling emergency medical service volunteer response," *Available at SSRN 3825060*, 2021.

[12] "Hartslagnu." (May 2022), [Online]. Available: `https://hartslagnu.nl/`.

[13] G. Slaa, "Increasing cardiac arrest survival by improving the volunteer alerting algorithm," M.S. thesis, University of Twente, 2020.

[14] D. Luxen and C. Vetter, "Real-time routing with openstreetmap data," in *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, ser. GIS '11, Chicago, Illinois: ACM, 2011, pp. 513–516, ISBN: 978-1-4503-1031-4. DOI: `10.1145/2093973.2094062`. [Online]. Available: `http://doi.acm.org/10.1145/2093973.2094062`.

[15] T. W. Schaap, L. W. J. Harms, M. Kansen, and H. Wüst, *Fietsen en lopen: de smeerolie van onze mobiliteit*. Kennisinstituut voor Mobiliteitsbeleid— KiM, 2015.

A Appendix

(a) *Incident type "Brand"*

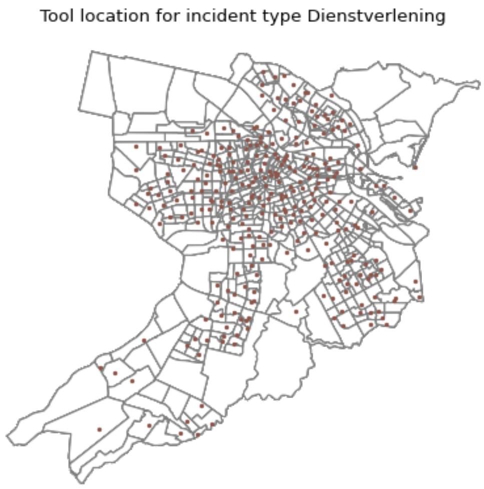

(b) *Incident type "Dienstverlening"*

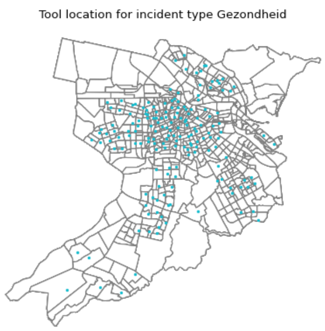

(c) *Incident type "Gezondheid" (healthcare)*

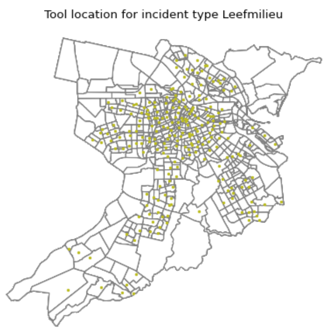

(d) *Incident type "Leefmilieu" (enviroment)*

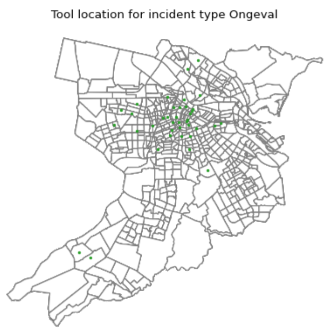

(e) *Incident type "Ongeval" (accident)*

Figure A.1: *Maps of tool locations for different incident types*

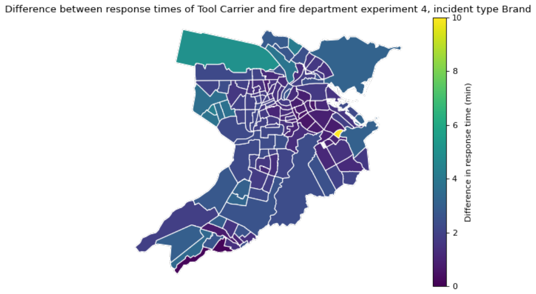

(a) *Incident type "Brand"*

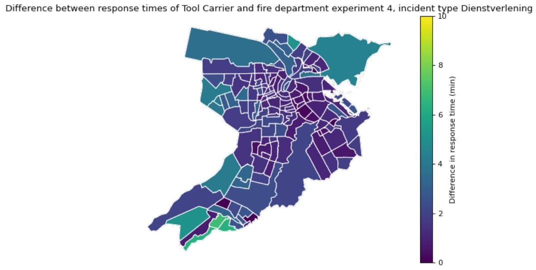

(b) *Incident type "Dienstverlening"*

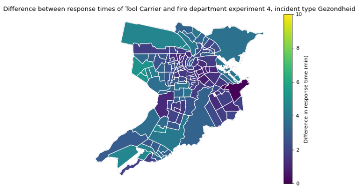

(c) *Incident type "Gezondheid" (healthcare)*

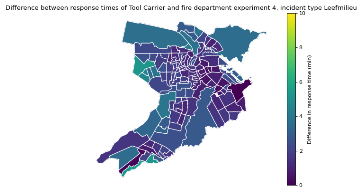

(d) *Incident type "Leefmilieu" (enviroment)*

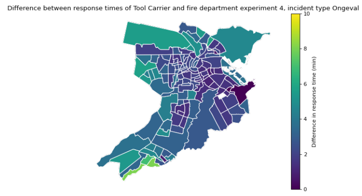

(e) *Incident type "Ongeval" (accident)*

Figure A.2: *Mean time gains for the different incident types of experiment 4 tool placement*

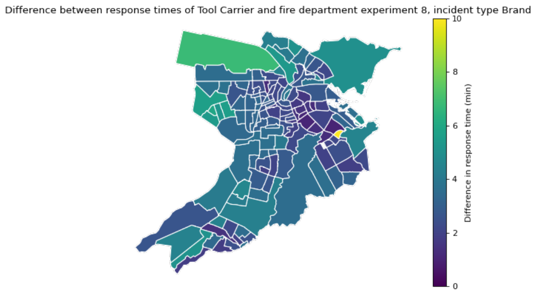

(a) *Incident type "Brand"*

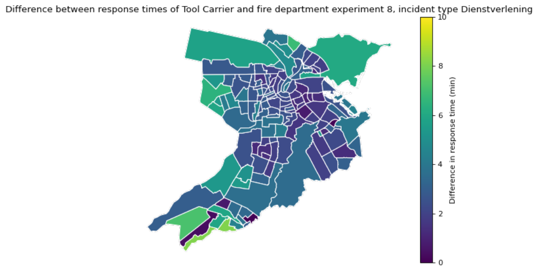

(b) *Incident type "Dienstverlening"*

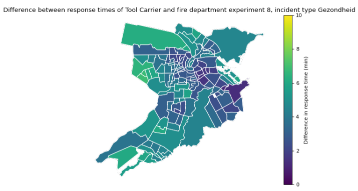

(c) *Incident type "Gezondheid" (healthcare)*

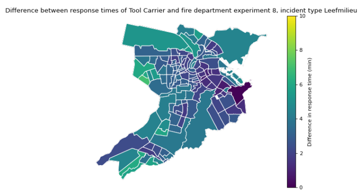

(d) *Incident type "Leefmilieu" (enviroment)*

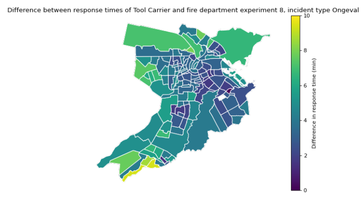

(e) *Incident type "Ongeval" (accident)*

Figure A.3: *Mean time gains for the different incident types of experiment 8 tool placement*

	Shapiro-Wilk test		Wilcoxon test	
Experiment pair	Statistic	p-value	Statistic	p-value
2 and 1 RR	0.98	0.0	2851465165	0.0
1 and 3 RR	0.97	0.0	701828904	0.0
1 and 3 TC	0.93	0.0	822597591	0.0
2 and 7 RR	0.91	0.0	469757417	0.0
4 and 8 TC	0.92	0.0	981763742	0.0
5 and 8 expert	0.25	0.0	3780590472	0.0

Table A.1: *Detailed results of pairwise testing*

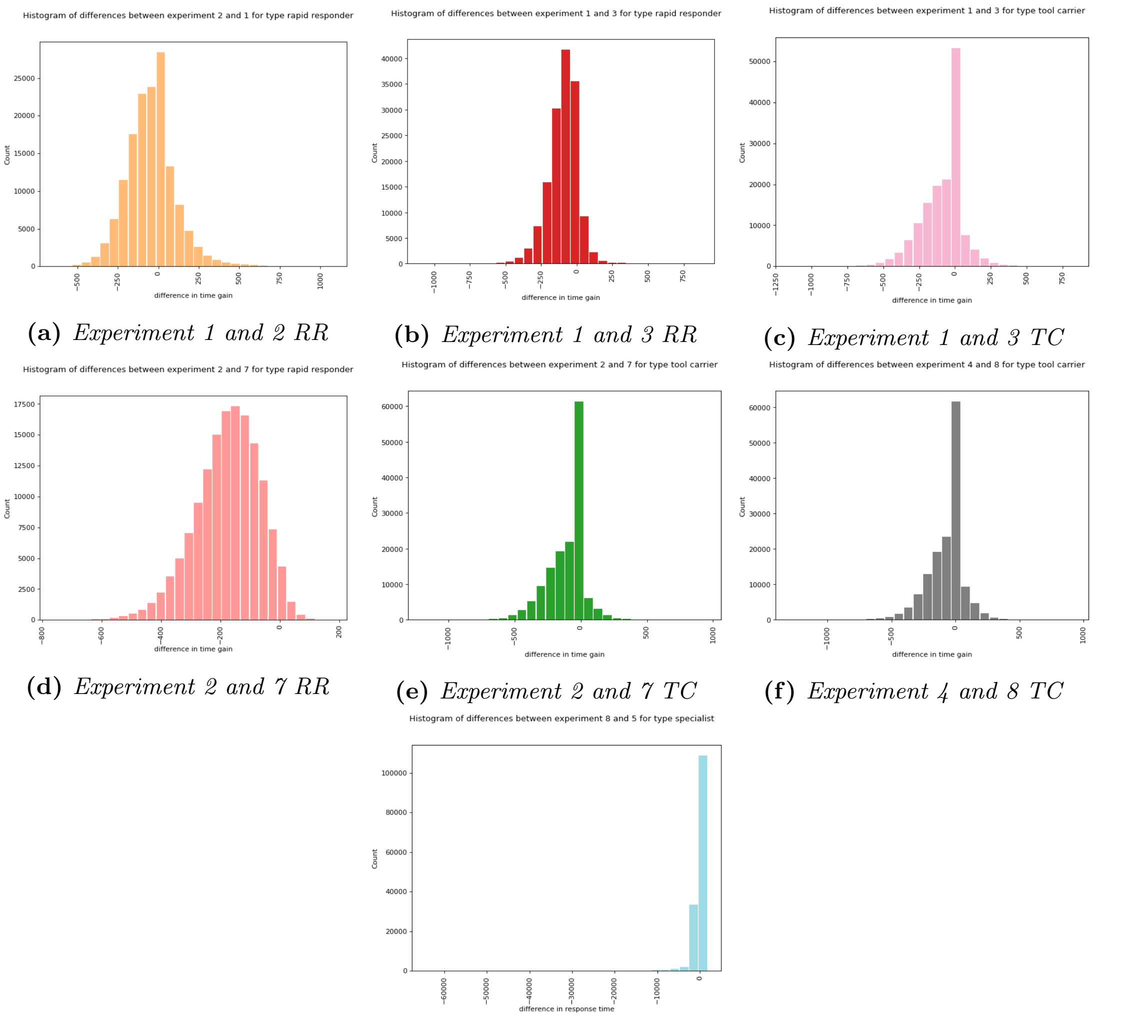

(a) *Experiment 1 and 2 RR*

(b) *Experiment 1 and 3 RR*

(c) *Experiment 1 and 3 TC*

(d) *Experiment 2 and 7 RR*

(e) *Experiment 2 and 7 TC*

(f) *Experiment 4 and 8 TC*

(g) *Experiment 5 and 8 specialist (response time, not time gain)*

Figure A.4: *Histograms of differences between samples of time gain of experiment pairs, to verify normality*