



Optimal patient protocols in regional acute stroke care

B. L. Garcia¹ · R. Bekker² · R. D. van der Mei³ · N. H. Chavannes¹ · N. D. Kruyt^{4,5}

Received: 25 September 2019 / Accepted: 4 October 2020
© The Author(s) 2021

Abstract

In acute stroke care two proven reperfusion treatments exist: (1) a blood thinner and (2) an interventional procedure. The interventional procedure can only be given in a stroke centre with specialized facilities. Rapid initiation of either is key to improving the functional outcome (often emphasized by the common phrase in acute stroke care “time=brain”). Delays between the moment the ambulance is called and the initiation of one or both reperfusion treatment(s) should therefore be as short as possible. The speed of the process strongly depends on five factors: patient location, regional patient allocation by emergency medical services (EMS), travel times of EMS, treatment locations, and in-hospital delays. Regional patient allocation by EMS and treatment locations are sub-optimally configured in daily practice. Our aim is to construct a mathematical model for the joint decision of treatment locations and allocation of acute stroke patients in a region, such that the time until treatment is minimized. We describe acute stroke care as a multi-flow two-level hierarchical facility location problem and the model is formulated as a mixed integer linear program. The objective of the model is the minimization of the total time until treatment in a region and it incorporates volume-dependent in-hospital delays. The resulting model is used to gain insight in the performance of practically oriented patient allocation protocols, used by EMS. We observe that the protocol of directly driving to the nearest stroke centre with special facilities (i.e., the mothership protocol) performs closest to optimal, with an average total time delay that is 3.9% above optimal. Driving to the nearest regional stroke centre (i.e., the drip-and-ship protocol) is on average 8.6% worse than optimal. However, drip-and-ship performs better than the mothership protocol in rural areas and when a small fraction of the population (at most 30%) requires the second procedure, assuming sufficient patient volumes per stroke centre. In the experiments, the time until treatment using the optimal model is reduced by at most 18.9 minutes per treated patient. In economical terms, assuming 150 interventional procedures per year, the value of medical intervention in acute stroke can be improved upon up to € 1,800,000 per year.

Keywords Operations research · Facility location · Allocation protocol · Mothership · Drip-and-ship · Mixed integer linear programming · Volume-dependent in-hospital delays

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: This work was supported by Health-Holland, Stichting LSH-TKI (project number LSHMI6041).

✉ B. L. Garcia
bjartygarcia@hotmail.com

- ¹ Public Health and Primary Care, Leiden University Medical Center, Leiden, Netherlands
- ² Department of Mathematics, Vrije Universiteit Amsterdam, Amsterdam, Netherlands
- ³ Stochastics Group, CWI, Amsterdam, Netherlands
- ⁴ Neurology Department, Leiden University Medical Center, Leiden, Netherlands
- ⁵ University NeuroVascular Center (UNVC), Leiden-The Hague, The Netherlands

Highlights

- Regional acute stroke care with two reperfusion treatments is modeled using a mathematical optimization model.
- The time until treatment is minimized based on the joint decisions of treatment locations and the allocation of acute stroke patients.
- The optimization model incorporates the impact of patient volume on in-hospital time delays.
- Experiments with six regions in the Netherlands showed that mothership is in most cases preferred over drip-and-ship and time until treatment can be reduced by at most 18.9 minutes per treated patient.

1 Introduction

Fast treatment of acute stroke is paramount and increases the prospects of good clinical outcome [1–8]. The passing of time is the most pivotal factor limiting clinical efficacy. For example, for every 15 minute delay to treatment, the chances of independent ambulation substantially decreases and 1 month of disability-free life is lost [9]. Two options are available to treat acute stroke, depending on patient characteristics. Both treatments aim to re-establish blood flow to the brain. The first is intravenous thrombolysis treatment (IVT) for which the majority of patients are eligible, but IVT is only effective within 4.5 hours from onset of stroke. The second is intra-arterial thrombectomy (IAT), which is a relatively new treatment option only applicable to a subset of selected ischemic stroke patients (4.5% according to Dutch Acute Stroke Audit), but can be applied up to 24 hours after symptom start. Importantly, in clinical practice, IAT is administered on top of IVT in approximately 21% of the IVT treated patients, but this percentage is increasing [19]. Moreover, of IAT treated patients, over 90% were first treated with IVT [25].

Whereas IVT is a treatment option available in most hospitals, so-called primary stroke centres (PSC's), IAT is only available in so-called comprehensive stroke centres (CSC's). In general, most regions possess a plurality of PSCs, but only a few CSCs, since the latter requires elaborate treatment techniques and specialized personnel. Patient allocation by emergency medical services (EMS) would ideally be based on IAT eligibility, but this requires in-hospital assessment. This means IAT eligibility can only be assessed at a centre and not at the patient location. It is therefore key to determine appropriate CSC locations on the regional level (i.e., decide which stroke centres are going to administer IAT for their area).

The following five factors influence the regional time delays in acute stroke care: (1) patient location; (2) regional patient allocation (i.e., regional allocation protocol prescribing to which type of stroke centre a suspected stroke patient is allocated by the EMS); (3) travel times of EMS; (4) PSC and CSC locations; and (5) in-hospital delays. It is not possible to influence the patient location, but the other factors can be controlled to some extent.

The focus of this paper is on the second and fourth factor, respectively, the patient allocation protocol and the PSC/CSC locations. These two factors also influence the other two controllable factors, namely the EMS travel times and in-hospital delays. Increased patient load in a stroke centre improves in-hospital logistics and typically leads to shorter delays in stroke care [14, 29–31]. In-hospital delays can thus be controlled to an extent through the clever allocation of patients. Currently, unsubstantiated decisions regarding regional patient allocation and PSC

and CSC locations are prevalent. A model-based method of organizing acute stroke care therefore offers great opportunities for reducing time delays until start treatment.

In clinical practice, two different patient allocation protocols are in use: *drip-and-ship* and *mothership* [17, 18, 28]. In the *drip-and-ship* protocol, suspected stroke patients are always allocated to the nearest stroke centre (either PSC or CSC) in order to initiate IVT as soon as possible. If a patient admitted to a PSC appears IAT eligible, subsequent allocation to a CSC follows. In general, this protocol results in short delays before IVT, but longer delays to IAT due to inter-hospital transportation. In the *mothership* protocol, suspected stroke patients are allocated directly to a CSC, often bypassing a PSC. In general, this protocol results in short delays before IAT, but relatively longer delays before IVT due to longer transportation times between patient location and CSC, compared to the shorter allocation time to a PSC. Essentially, neither protocol is optimal for a given region. The optimal protocol for a certain region may show elements of both the *mothership* and the *drip-and-ship* protocols, but may also use another allocation (e.g. allocate some demand to a PSC on the route to a CSC). The optimal allocation follows from the optimization model as introduced in Section 4.

Next to the allocation protocol, a strategic decision is the determination of which stroke centre should only administer IVT (i.e. become a PSC) and which should also administer IAT (i.e. become a CSC). To guarantee a certain degree of expertise at a CSC, a lower bound on the number of acute stroke patients offered at a location is required. In addition, the volume of acute stroke patients affects in-hospital delay [14, 29–31]. As IAT becomes a more familiar treatment option, this is an opportune moment to make well-founded strategic location decisions for PSCs and CSCs.

In the current paper, we propose a mathematical optimization model upon which to base the allocation protocol and PSC and CSC locations to minimize regional time delays, taking the impact of patient volume on in-hospital delays into account. The model is formulated as a mixed integer linear program (MILP) and solved using CPLEX. We also compare our optimized patient allocation protocol with *drip-and-ship* and *mothership* protocols, providing insight in the differences in performance. This research focuses on reducing delays to treatment initiation for the patient.

We see that the relative performance of *mothership* and *drip-and-ship* strongly depends on the region (population densities, stroke centre locations) as well as treatment-related parameters. As a rule of thumb, we observe the following: *mothership* performs well in urban areas and in regions where the CSCs are centrally located. Also, we see that *mothership* typically shows near-optimal performance when the fraction of IAT eligibility exceeds 50%. Vice

versa, drip-and-ship performs well in rural areas and fractions of IAT eligibility of 30% or smaller, assuming that the patient volume per stroke centre is sufficiently large.

There are two related areas of existing literature. The first area stems from the clinical domain for strokes. For instance, [16, 20] assess eligibility of IAT through clinical trials. More closely related are [15] and [17]. The paper [15] considers some scenarios for patient allocation based on clinical data, whereas [17] analyzes the difference between the drip-and-ship and mothership protocols based on current locations and allocation characteristics.

The second area of related work concerns the area of facility locations problems. There is a long history on location problems and a sizeable amount of literature by now. We refer to the book [40] for an overview of location problems in general, including some optimization models. The book contains a chapter dedicated to location problems in health care [39]; we also refer to [12] for a survey on health care facility location problems. In the drip-and-ship protocol, patient are allocated to the nearest stroke centre. The studies [41, 42] also consider health care settings that incorporate closest assignment constraints in the optimization model.

Another branch of literature takes congestion into account for facility location problems and the allocation of clients to resources. The allocation prescribes the arrival rates at locations, giving rise to queueing phenomena [43]. Such models lead to non-linear optimization models, which are optimized using heuristics in [44, 45]. In [46, 47] the goal is to minimize travel times plus waiting times in a health care network, where the waiting times follow from queueing expressions; the authors exploit an MILP combined with piece-wise linear approximations to linearize the goal function and all constraints.

Due to the two types of treatments and centres, the problem falls in the class of multi-flow two-level hierarchical facility location problems [10, 11]. An example of a two-level hierarchical model with a proportion of patients being transferred between health centres and hospitals is [48]; the optimization model is solved using five heuristic procedures. In [49] an MILP model is used for location-allocation problems with district and central hospitals. The design of a long-term multi-facility care network using large scale MILP's is considered in [34, 51]. Also of interest is [50], which considers a location-allocation problem that aims to optimize health gains directly using an MILP.

Even though many researches modelled healthcare networks using MILP, our understanding is that there are no studies that incorporate acute patients in a chain of stroke centres and volume-dependent delays in the facility location literature.

Our contribution is two-fold. First, we formulate optimization models for locations and allocation protocols for treatment of acute stroke in the spirit of two-level facility location models. Specifically, we incorporate volume-dependent in-hospital delays and practically common protocols (drip-and-ship and mothership). Second, we provide insight in the performance of the practically oriented protocols. The optimization framework is generic, but we imagine that the practical heuristics may be preferred in practice.

The paper is organized as follows. In Section 2 we give some background and introduce notation for the modeling of stroke care. Some initial insights for the different protocols is presented in Section 3 based on simulation results of a stylized example. The optimization formulation is given in Section 4. Numerical experiments can be found in Section 5, whereas Section 6 concludes.

2 Modeling regional acute stroke care logistics

In this section, we provide some background about acute stroke care and introduce notation for the model. In Fig. 1, the possible patient flows from onset to start treatment are depicted. We focus on the flow starting from departure ambulance from patient location and ending with the start of either IVT or IAT, as depicted by the red area in Fig. 1. We call this time the SDST (Scene-Departure-to-Start-Treatment). Thus, the time from onset until the ambulance arrives at the scene is left out of scope; this duration depends on many external factors and is not affected by patient allocation protocols and PSC and CSC locations. Below, we introduce the elements of the model.

Location A region is assumed to consist of a finite number of non-overlapping areas, called locations. Typically, a location corresponds to a postal code area (as in e.g. [13]). The non-overlapping areas (or locations) are relatively small such that the distance within a location may be assumed to be zero. We assume that demand may occur from each location, whereas there is a smaller set of treatment locations of stroke centres.

Demand For the demand we consider all IVT- and IAT-eligible patients. The suspected stroke patients that do not require treatment are not affected by the allocation protocol and the PSC and CSC locations. Therefore, this group is considered to be out of scope. Note that this group does affect the occupation of EMS to some extent. We denote the demand per year at demand location i by w_i , and let $P = \sum_i w_i$ be the total demand in the region per year.

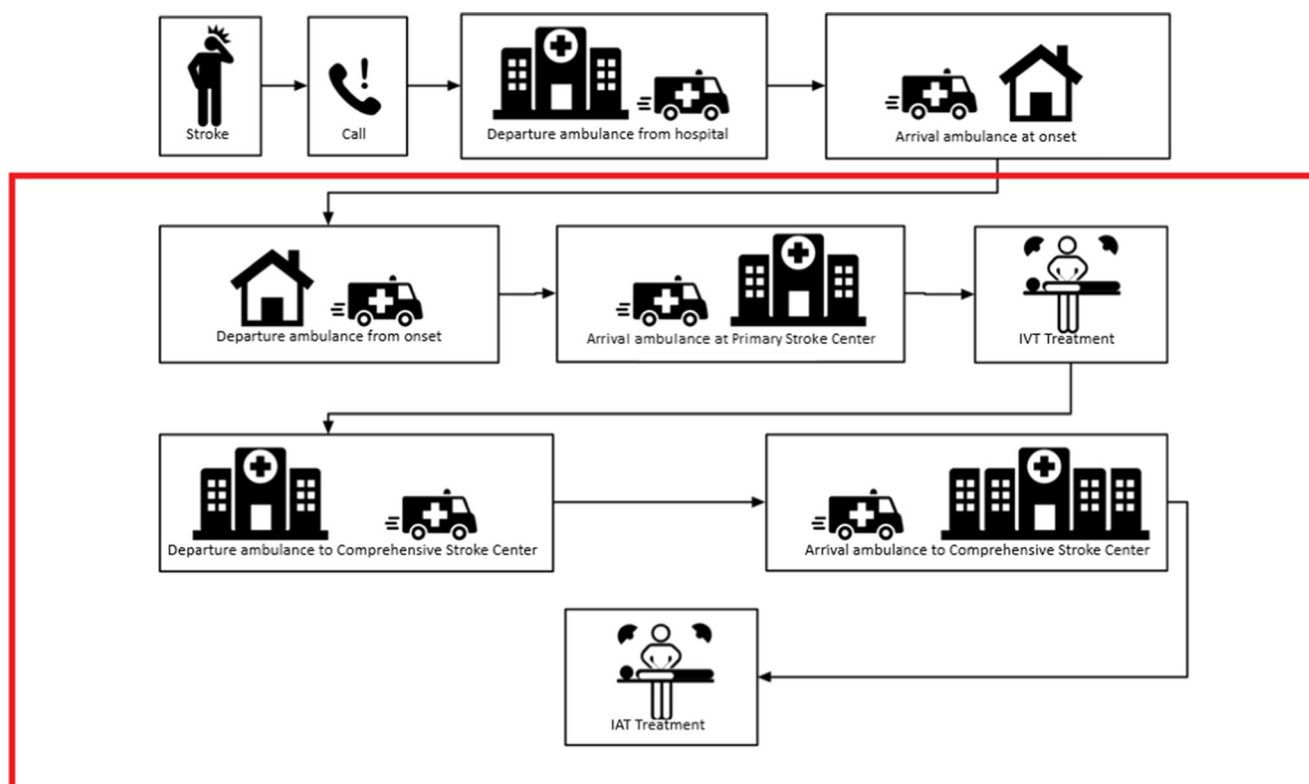


Fig. 1 Visualization of the patient flow for acute stroke; we focus on the boxed area

Typically, the demand per location depends on demographic characteristics and can be retrieved from open sources.

PSC and CSC locations The model allocates treatments (IVT and/or IAT) to the potential treatment locations. These typically correspond to existing stroke centres and are a subset of all locations in a region. Treatment locations may correspond to the locations of demand of IVT- and IAT-eligible patients, but may also be considered as separate locations.

Travel time We denote the travel time between locations i and j by d_{ij} . In this study, travel times are based on average emergency travel time estimates between two locations provided by the Dutch National Institute for Public Health and the Environment (RIVM). The model can also be used with travel times from applications (e.g. Google Maps) or distance measures; for the latter, we refer to [21] for a more elaborate discussion.

Eligibility of IAT treatment The probability of being IAT eligible, denoted by $P(IAT)$, depends on individual patient characteristics [22], but not on location or transportation. In this model, $P(IAT)$ is constant over time and is defined as the fraction of IVT patients that are eligible to receive IAT.

In-hospital delays The in-hospital delay is the time from arrival at the door of the stroke centre until the start of the specific treatment (IVT or IAT). In-hospital delays inversely depend on the volume of IVT treated patients with higher volumes leading to shorter delays [14, 29–31]; larger patient volumes provide experience and often also lead to organizational improvements. In addition, very small patient volumes are undesirable; a sufficient volume of patients should be treated to ensure that skills are maintained and efficiency is not becoming an issue [32]. Hence, we assume that in-hospital delays behave as a convex function of the patient volume.

Minimum volume requirement for IVT/IAT To maintain a sufficient level of expertise in the treatment of acute stroke, it is crucial to treat a minimum number of patients. In clinical practice, there is a minimal requirement for the volume of IVT/IAT patients [26]. In our model, the minimal number of IVT and IAT patients per year are denoted by r_{IVT} and r_{IAT} , respectively.

3 Insights from a stylized example

In this section, we investigate the impact of the different allocation protocols that are currently in use by considering

a stylized example. The primary goal is to obtain insight in how the allocation protocols - mothership, drip-and-ship, and optimal - behave for different values of $p := P(IAT)$ and the number of PSC facilities relative to the number of CSCs. We assume that the region of interest is a square and denote this square by U ; without loss of generality we assume U to be the unit square. There is a single CSC and there are n PSCs, which are uniformly distributed over U . The demand point is also uniformly distributed over U . The distance d_{ij} between two points i and j is given by the Euclidean distance (the L^2 norm); see Remark 3 for the Manhattan distance.

Note that if the location of facilities are given and in-hospital delays are independent of the number of patients treated at the facilities (and there are no upper and lower bounds), the optimal allocation protocol can easily be obtained. This is due to the fact that the optimization can be separately carried out for each demand point. The optimal allocation for demand point i is then determined by

$$\arg \min_{k=0, \dots, n} \{d_{ijk} + b_{k,IVT} + p(d_{jkj_0} + b_{IAT})\}, \tag{1}$$

where j_0 is the CSC location and j_k is the location of PSC $k, k = 1, \dots, n$ (with $d_{j_0, j_0} = 0$). Here, $b_{k,i}$ denotes the in-hospital delay for treatment $i \in \{IVT, IAT\}$ at location k . For mothership, the SDST is $d_{ij_0} + b_{0,IVT} + pb_{IAT}$, whereas for drip-and-ship the SDST is

$\min_{k=0, \dots, n} \{d_{ijk}\} + b_{k^*,IVT} + p(d_{j_k^*j_0} + b_{IAT})$, where $k^* = \operatorname{argmin}_{k=0, \dots, n} d_{ijk}$. In this example, we focus on the travel time and exclude in-hospital delays (that is, assume them to be identical).

In Fig. 2, the mean of the travel distances are shown for the three different protocols and for $n = 1$ and $n = 4$ based on 10,000 simulations per instance. We varied p to see the impact of the three protocols. Naturally, the mothership protocol does not depend on p and n . The performance (generally) improves with 4 PSCs compared to the situation with 1 PSC, which may be evidently explained




Fig. 2 Mean of the travel distances for various $p = P(IAT)$ and allocation protocols

by a smaller distance to the nearest facility and more options for routing patients in case of the optimal protocol. Moreover, the optimal protocol clearly outperforms the other two protocols (for the same number of PSCs). When $p \downarrow 0$ drip-and-ship is optimal as there is no need for further allocation of patients to a CSC. On the other hand, for $p \uparrow 1$, mothership is optimal, as all patients should eventually be allocated to a CSC. For $p = 0.5$ the relative gain in mean travel distance of the optimal protocol compared to mothership is roughly 7% and 18% for $n = 1$ and $n = 4$, respectively.

Comparing mothership with drip-and-ship, we see that drip-and-ship performs better for p smaller than 0.5 in case $n = 1$, where roughly $p = 0.5$ is the break-even point. For $n = 4$, the break-even point shifts to a slightly larger value of p . Interestingly, when p is above 0.8 the drip-and-ship protocol with $n = 1$ has a smaller mean distance than for the case that $n = 4$. This is due to the fact that it is more likely that the CSC is the nearest location in case $n = 1$ than for $n = 4$.

The example above provides some fundamental insight in how $P(IAT)$ and the number of PSCs affect the allocation protocols. Moreover, Equation (1) provides a simple heuristic for determining allocation strategies. With volume-dependent IVT treatment times and different demographic regions, the patterns described in this section will roughly remain valid, but the performance loss may shift. In Section 5, we elaborate on this when demographic regions are taken into account in an optimization model.

Remark 1 Distances between random points have been extensively studied in the literature on spatial point processes. For instance, the average distance between two random points in the unit square is

$(2 + \sqrt{2} + 5(\ln(1 + \sqrt{2}))) / 15 \approx 0.5214$, see p. 171 of [23], corresponding to the simulation results for mothership. We also refer to [23, 24, 27] for further examples. Observe, however, that the drip-and-ship and optimal allocation protocols are considerably more intricate to derive analytically.

Remark 2 It may be easily verified that there is a p^* such that drip-and-ship outperforms mothership for $p < p^*$, and vice versa for $p > p^*$. This follows from the intermediate value theorem, the fact that the performance of drip-and-ship is decreasing in p (and independent of p for mothership), and the boundary cases $p = 0$ and $p = 1$.

Remark 3 Next to the Euclidean distance, another common distance metric is the Manhattan distance (L^1 norm). The distance between points (x_1, y_1) and (x_2, y_2) is then $|x_1 - x_2| + |y_1 - y_2|$. Although the actual distances for the L^1

norm differ from those of the L^2 norm, the relative results are very similar and are thus not presented here.

4 Optimization framework

In this section, we propose an optimization model for jointly determining the allocation protocols and PSC and CSC locations, in order to minimize the mean SDST. The model is formulated as an MILP based on a modification of a multi-flow two-level hierarchical facility location problem [11]. A primary feature that we include is that in-hospital delays depend on patient volumes. In addition to the optimal allocation protocol, we also discuss the model for optimal locations in case drip-and-ship is enforced.

4.1 Parameters and variables

An outline of regional stroke care, including notation, has been given in Section 2. As mentioned, we assume that there is a minimum requirement r_c for patients treated of type $c \in \{IVT, IAT\}$ at open stroke centres. Especially because IAT is still relatively new, we also include a maximum number of PSC and/or CSC locations in the region, denoted by p_c . The binary decision variable h_{jc} equals 1 if stroke centre at location $j \in J$ provides treatment type $c \in C$, and 0 otherwise.

The allocation of patients from their demand location to the PSC is given by the binary decision variables y_{ij} ($i \in I, j \in J$), where y_{ij} takes value 1 if a patient from demand point i is allocated to stroke centre location j , and 0 otherwise. Subsequently, the potential allocation of patients from a PSC to a CSC is determined by v_{ij} , which equals the total number of patients from PSC location $i \in J$ to CSC location $j \in J$. These locations can be identical in case a stroke centre provides both IVT and IAT. The distance d_{jj} ($j \in J$), may take negative values to model the situation where there is an efficiency gain in delays when the patient remains in the same stroke centre. Thus, $d_{ij} \geq 0$ for ($i \in I \cup J, j \in J$) and $j \neq i$, whereas $d_{jj} \in \mathbb{R}$.

We assume that in-hospital delays are stroke centre specific, i.e., let $b_{j,c}$ be the in-hospital delay for location j and treatment of type c . As a relation between the number of IAT patients and in-hospital delays has not yet been established, we assume this delay to be independent of the volume. From [14], it follows that in-hospital delay of IVT depends on volume according to a convex function. Specifically, $b_{j,IVT}$ is a function of the volume of the location $v := \sum_i w_i y_{ij}$; that is the IVT time $b_{j,IVT} = b(v)$ is convex in the volume v . In our MILP model, we approximate the in-hospital delay of IVT by using N piece-wise linear functions, see Fig. 3 for an example with $N = 3$. These linear functions may be constructed from $b(v)$ based

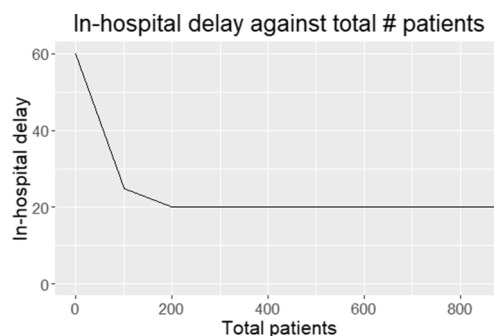


Fig. 3 Example of the piece-wise linear approximation of the in-hospital delay as a function of volume of IVT treatments

on the tangent lines of $b(v)$ in the points $v_n, n = 1, \dots, N$. The n th tangent line $\hat{b}_n(\cdot)$ has the form $\hat{b}_n = \alpha_n \cdot v + \beta_n$, where the slope α_n and intercept β_n of the n -th linear function should be chosen such that $b(v_n) = \hat{b}_n(v_n)$ and $b'(v_n) = \alpha_n$ (Table 1).

As $b_{j,IVT}$ depends on the volume and thus on the allocation protocol, we need an auxiliary variable z_{ij} representing the total in-hospital delay for IVT for patients from demand location i receiving IVT at PSC location j . We refer to Table 2 for notation related to the IVT in-hospital delay and to Table 1 for the rest of the notation.

Remark 4 For acute stroke, we are primarily interested in scenario's where in-hospital delays are a decreasing function of patient volume. In case patient volumes form a considerable part of the overall volume offered at an emergency department, large volumes may also lead to congestion, which may be characterized using queueing models [43–47]. It is well known that the mean waiting times of the traditional M/M/c and M/G/1 queues are a convex function of the offered load and thus the arrival rate (i.e. volume of patients). As such, the model design is also applicable to situations where congestion phenomena occur; this is not the case for stroke care in our setting as the average number of newly arriving stroke patients is at most two per day.

4.2 MILP formulation

We now formulate the MILP for the optimal PSC and CSC locations and allocation protocol. The solution of the MILP provides the desired PSC and CSC locations ($h_{j,IVT}$ and $h_{j,IAT}$) and the corresponding optimal allocation protocol (y_{ij} and v_{ij}). The objective is to minimize the sum of the SDST (Scene-Departure-to-Start-Treatment) over all patients in the region. The objective function, i.e., the total SDST, in (2) is divided in two parts. The first part is the total delay from scene departure until IVT. The second part is the total additional delay from IVT until start of IAT.

Table 1 Notation for optimization models

Sets	
I	Set of demand locations (of suspected stroke).
J	Set of potential stroke centre locations.
C	Set of treatments types ($C = \{IVT, IAT\}$).
Parameters	
$P(IAT)$	Probability that patient requires IAT treatment.
w_i	Demand at location i .
d_{ij}	Travel time from location i to location j .
r_c	Minimum number of patients with treatment c .
p_c	Maximum number of stroke centres that provide treatment type c .
$b_{j,IAT}$	In-hospital delay before start of IAT for stroke centre j .
Decision variables	
y_{ij}	1 if demand point i is assigned to PSC j for IVT, 0 otherwise.
v_{ij}	Flow of patients from PSC i to CSC j for IAT at location j .
h_{jc}	1 if treatment of type c is given at location j , 0 otherwise.

Now, let M be a sufficiently large number. The constraints (3) make sure that every demand point is assigned to a PSC. Constraints (4) regulate the allocation of patients from a PSC to a CSC in case IAT is required. In particular, it states that the flow of patients into PSC j that need IAT equals the flow out of PSC j . Constraints (5) and (6) guarantee that the minimum number of IVT and IAT patients are treated in case the location is open as a PSC and CSC, respectively. Similarly, constraints (7) and (8) provide that stroke centre j needs to be open for IVT and IAT, respectively, in case IVT and IAT patients are assigned to stroke centre j . Constraints (9) makes sure that the number of stroke centres of type c is less than or equal to p_c .

$$\min Z_{opt} = \sum_{i=1}^{|I|} \sum_{j=1}^{|J|} [y_{ij} w_i d_{ij} + z_{ij}] + \sum_{j=1}^{|J|} \sum_{c=1}^{|C|} [v_{jc} (d_{jc} + b_{c,IAT})] \tag{2}$$

$$\text{s.t.} \quad \sum_{j=1}^{|J|} y_{ij} = 1 \quad \forall i \tag{3}$$

$$\sum_{i=1}^{|I|} y_{ij} w_i P(IAT) = \sum_{c=1}^{|C|} v_{jc} \quad \forall j \tag{4}$$

Table 2 Notation for piece-wise linear approximation of IVT in-hospital delay

α_n	Slope of the n th tangent line of IVT in-hospital delay.
β_n	Intersect of the n th tangent line of IVT in-hospital delay.
$b_{j,IVT}$	In-hospital delay for IVT at stroke centre j .
z_{ij}	Total in-hospital delay for IVT for demand i in stroke centre j .

$$\sum_{i=1}^{|I|} y_{ij} w_i \geq r_{IVT} h_{j,IVT} \quad \forall j \tag{5}$$

$$\sum_{i=1}^{|I|} v_{ij} \geq r_{IAT} h_{j,IAT} \quad \forall j \tag{6}$$

$$\sum_{i=1}^{|I|} y_{ij} \leq M h_{j,IVT} \quad \forall j \tag{7}$$

$$\sum_{i=1}^{|I|} v_{ij} \leq M h_{j,IAT} \quad \forall j \tag{8}$$

$$\sum_{j=1}^{|J|} h_{jc} \leq p_c \quad \forall c \tag{9}$$

$$z_{ij} \leq M y_{ij} \quad \forall i, j \tag{10}$$

$$z_{ij} \leq b_{j,IVT} w_i \quad \forall i, j \tag{11}$$

$$z_{ij} \geq b_{j,IVT} w_i - (1 - y_{ij}) M \quad \forall i, j \tag{12}$$

$$b_{j,IVT} \geq \sum_{i=1}^{|I|} \alpha_n w_i y_{ij} + \beta_n \quad \forall n, i, j \tag{13}$$

$$y_{ij} \in \{0, 1\} \quad \forall i, j \tag{14}$$

$$v_{jj'} \geq 0 \quad \forall j, j' \tag{15}$$

$$h_{jc} \in \{0, 1\} \quad \forall j, c \tag{16}$$

$$z_{ij} \geq 0 \quad \forall i, j \tag{17}$$

$$b_{j,IVT} \geq 0 \quad \forall j \tag{18}$$

Observe that the total in-hospital delay from the demand at location j is $\sum_i y_{ij} w_i b_{j,IVT}$, where the in-hospital delay per patient $b_{j,IVT}$ depends on IVT volume at the corresponding location. Equations (10)–(12) avoid the multiplication of the variables y_{ij} and $b_{j,IVT}$ by introducing

the auxiliary variable z_{ij} . Specifically, the equations make sure that $z_{ij} = 0$ if $y_{ij} = 0$ and $z_{ij} = w_i b_{j,IVT}$ if $y_{ij} = 1$. Constraints (13) model the in-hospital delay per patient at location j depending on total IVT volume at location j using its convex relationship (and approximate this function by n piece-wise linear functions). Finally, the constraints (14)–(18) guarantee that all decision variables are non-negative and that y_{ij} and h_{jc} are binary.

4.3 Drip-and-ship and mothership

Although the allocation protocol is given in this section, an optimization model can be used to determine which treatment locations should be a CSC (and then also a PSC for drip-and-ship). In fact, using the optimization model it may be determined whether performance is improved when some stroke centres no longer act as a PSC and CSC. For the results, we do not take such scenario's into account. Next, we first discuss the drip-and-ship protocol, followed by some comments regarding mothership.

The difference between the drip-and-ship protocol and the formulation in Section 4.2 is that we need to ensure that a patient is assigned to the nearest PSC. For any PSC i , let J_{ij}^* be the set of PSCs that are closer to demand i than PSC j (which may also be an empty set). Then, in addition to all constraints in Section 4.2, we need to impose the following constraints

$$\sum_{k \in J_{ij}^*} h_{k,IVT} \leq (1 - y_{ij})M \quad \forall i, j$$

This constraint regulates that no PSC in the set J_{ij}^* may be opened if demand i is assigned to PSC j , i.e., demand should be assigned to the nearest open PSC. M is again a sufficiently large number.

Observe that in the formulation above, some of the stroke centres may no longer provide IVT (and should then not be regarded as a stroke centre anymore). Although this improves the mean SDST, a typical practical implementation may be a scenario where all considered stroke centres should at least provide IVT. This can easily be accomplished in the model by modifying Constraints (9) into

$$\sum_{j=1}^{|J|} h_{jc} = p_c \quad \forall c \quad (19)$$

combined with the appropriate choice of p_c .

In case of the mothership protocol, the model reduces to a single-flow system as stroke centres can no longer be only PSCs. As such, the problem is actually a regular facility location problem. The key distinguishing feature, however, is that in-hospital delays still depend on patient volume. The formulation of the MILP for mothership can be found in Appendix A.

5 Numerical experiments

In this section, we compare the different allocation protocols using six different regions that differ in the number of demand points, population size, surface area, demand distribution and the potential stroke centre locations. In Section 5.1 the experimental setup is explained. In Section 5.2 a single instance is highlighted for illustration. Finally, in Section 5.3 a comparison of the different allocation protocols is given, with the computational time given in Section 5.4.

5.1 Regional structure and parameters

In this subsection, we first discuss the regional structures, followed by the choices of parameters.

We apply the models to 6 of the 24 ambulance regions (RAVs) in the Netherlands: Amsterdam-Amstelland, Holland-Midden, Haaglanden, Utrecht, Twente, and Groningen. The regions are chosen based on population density, following the classification in [38]; that is, we consider urban (Amsterdam & Haaglanden), rural (Groningen & Twente) and mixed (Holland-Midden & Utrecht) regions. As demand points, we take the four-digit postal codes.

Figure 10 in Appendix B shows the lay-out of the six regions. The surface area varies between 282 km^2 (Amsterdam-Amstelland) and 2336 km^2 (Groningen). Recall that the total demand in a region is equal to the model parameter P . We assume that every demand location is assigned a fraction of P proportional to the population density provided by the RIVM. For the potential PSC and CSC locations per region, we follow the categorization of [37]. In Fig. 10, the PSCs are indicated with a small circle and the CSCs with a cross.

Next to the regional structure, we use the parameters as indicated in Table 3. For each region, this provides 39 instances¹ with a different parameter combination. The probability of IAT eligibility is based on expert opinion about the range of this probability [20]; in current clinical practice 20% is realistic, but this value is expected to increase. The minimum IVT requirement (r_{IVT}) is set to '0' to mimic current practice and avoid that the drip-and-ship protocol becomes infeasible due to the patient volume for one of the PSCs dropping below r_{IVT} . Finally, the in-hospital delay for IVT is chosen according to the function in Fig. 3, ranging from 20 to 60 minutes, depending on patient volume. The in-hospital delay for IAT ($b_{j,IAT}$) is 29 minutes.

The allocation protocols are compared based on the total delay from scene departure to start treatment (i.e., the SDST), corresponding to the objective function Z of the

¹We only consider parameter combinations leading to feasible instances. For instance, with $P(IAT) = 20\%$ and $P = 300$, there are only 60 patients requiring IAT, such that $r_{IAT} > 60$ is not feasible. Consequently, we disregard six parameter combinations per region.

Table 3 Parameter values for numerical experiments

Parameter	Experiment values
$P(IAT)$	{20%, 30%, 40%, 50%, 60%}
Total number of patients (P)	{300, 600, 900}
Minimum IAT requirement (r_{IAT})	{50, 100, 150}

MILP. The relative difference in SDST for drip-and-ship and mothership, compared to the optimal protocol, is denoted by

$$\Delta_p = \frac{Z_p - Z_{opt}}{Z_{opt}} \times 100\%,$$

where Z_p is the objective function for protocol p and Z_{opt} is the objective function of the optimal model.

5.2 Insights from a single instance

To obtain insight in the impact of the different protocols, we first elaborate on a single instance. We focus on the Amsterdam-Amstelland region with $P(IAT) = 20\%$, the total number of patients $P = 600$ per year and a minimum IAT requirement r_{IAT} of 50 per CSC.

The results of this Amsterdam-Amstelland instance is given in Table 4. The optimal allocation is better than drip-and-ship by 7.7% and than mothership by 11.5%. Moreover, the number of PSCs differs between the protocols; the optimal protocol has one less open PSC than drip-and-ship, whereas there are by definition no open PSCs for mothership. Finally, the fraction of patients that require IAT and need to be transferred is 90.3% for the drip-and-ship protocol; only 9.7% of the IAT patients are directly allocated to a CSC. For the optimal protocol, the percentage of required IAT transfers decreases slightly to 81.0%. We note that these high number of IAT transfers can be explained by the fact that the CSC is on the outskirts of Amsterdam and at the edge of the region.

In Figs. 4, 5 and 6 the allocation of demand to the stroke centres is indicated by black arrows. Figure 6 shows how patients are allocated to the nearest PSC, including the PSC at the top of the region, as required by the drip-and-ship protocol. Figure 4 shows that in the optimal protocol the PSC at the top of the region is no longer used for IVT patients. Also, for many locations, the optimal protocol allocates patients to the nearest PSC. An exception is the

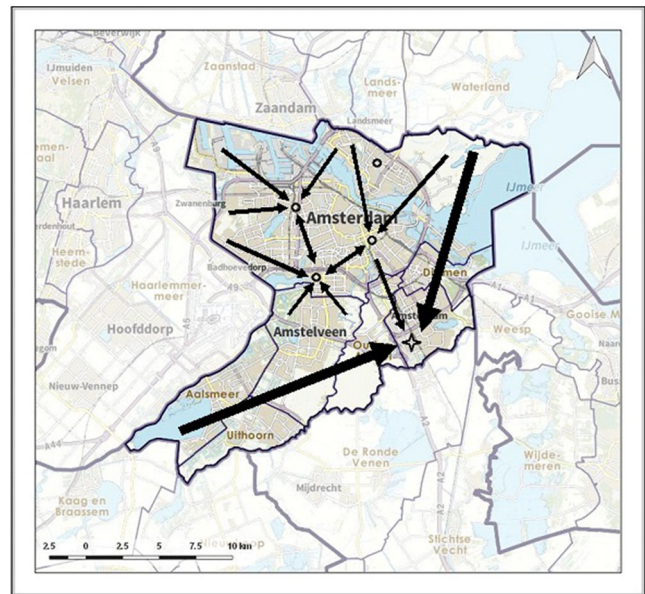


Fig. 4 Allocation of demand according to optimal model

southeastern area, for which it is better to allocate patients directly to the CSC, which is similar to the allocation in the mothership protocol (as illustrated in Fig. 5).

For this single instance we observe that drip-and-ship outperforms mothership. This can be explained by the small percentage of patients that require IAT and the location of the CSC. Of course, results will strongly depend on the regional layout and health-related parameters.

5.3 Analysis of different protocols

We compare the protocols for each region using the 39 feasible parameter combinations, resulting in a total of 234 instances.

Let us first focus on the difference between drip-and-ship and mothership. Table 5 shows the fraction of instances (in %) that mothership outperforms drip-and-ship in terms of SDST for the different regions and different values of $P(IAT)$. Not surprisingly, and in line with Section 3, we see that mothership will perform better compared to drip-and-ship as $P(IAT)$ increases. Next to $P(IAT)$, the trade-off between the two practical protocols strongly depends on the region and the locations of stroke centres. For instance, we see that for the region Haaglanden mothership

Table 4 Results for single ‘Amsterdam-Amstelland’ instance ($P(IAT) = 20\%$, $P = 600$, $r_{IAT} = 50$)

	Optimal	Mothership	Drip-and-ship
Sum of SDST	19069	21274	20531
Relative difference Δ	-	11.5%	7.7%
Number of (PSC, CSC)	(3,1)	(0,1)	(4,1)
Fraction of transferred IAT patients	81.0%	0%	90.3%

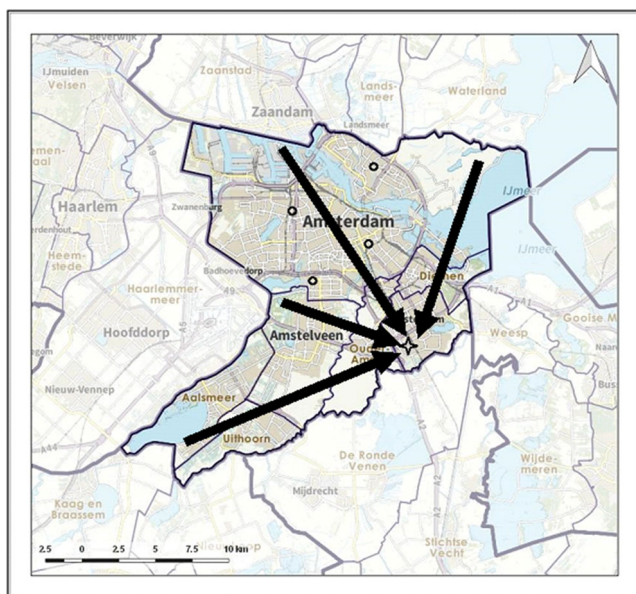


Fig. 5 Allocation of demand according to mothership

outperforms drip-and-ship for all instances. This can be explained by the degree of urbanization and corresponding short travel times. For the rural area Twente, we observe the opposite; only in some cases and for a $P(IAT)$ of at least 50% it holds that mothership has shorter SDST than drip-and-ship. Roughly speaking, we can see that mothership gives better performance for urban areas, whereas drip-and-ship works well for rural areas, which can be explained by the impact of travel times.

In addition, the location of the CSC(s) seem to play a crucial role. Amsterdam and Haaglanden are both urban areas, but for Amsterdam the CSC is located near the

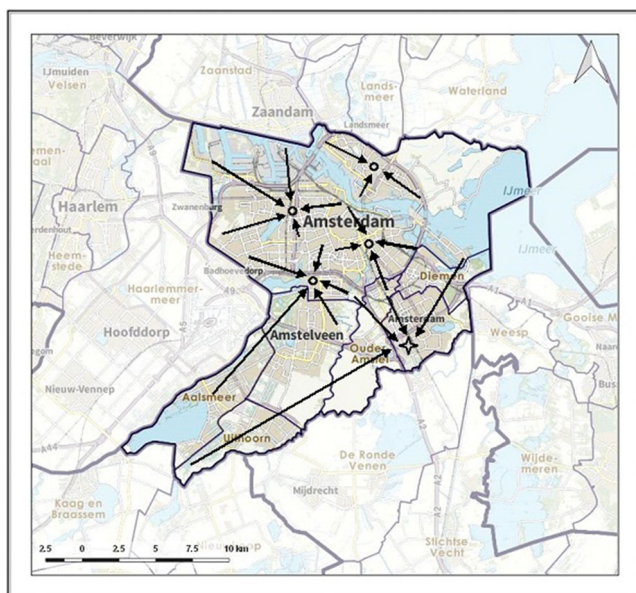


Fig. 6 Allocation of demand according to drip-and-ship

edge of the region. In Amsterdam, for smaller values of $P(IAT)$, the longer travel times to a CSC in the mothership protocol do not outweigh the additional travel times due to transfers of IAT patients in drip-and-ship. Similar arguments apply to the other regions; the CSCs in Twente (urban) and Holland-Midden (mixed) are closer to the edge of the region compared to Groningen (urban) and Utrecht (mixed).

In Fig. 7 the relative differences (Δ_p) between drip-and-ship (left) and mothership (right) compared to the optimal protocol are visualized using boxplots based on the 39 instances per region. The observations above concerning Table 5 also apply to a large extent to the performance per region in Fig. 7. For instance, we see again that overall mothership performs better than drip-and-ship in the urban regions, where the reverse holds for rural areas (except for some cases with $P(IAT) \geq 50\%$). The variability in performance for drip-and-ship can be considerably larger than for mothership. For example, for the regions Amsterdam and Utrecht the performance of drip-and-ship is almost 40% worse than optimal for $P(IAT) = 60\%$ and $P = 300$. For mothership, the performance is only roughly 12% worse than optimal for the regions Amsterdam, Holland-Midden, and Twente in case $P(IAT) = 20\%$ and $P \geq 600$. On average the total SDST when using the drip-and-ship protocol is 8.6% larger than the optimal model, whereas for the mothership protocol this is only 3.9%. Nevertheless, the mothership protocol can perform worse than the drip-and-ship protocol in all regions (except Haaglanden). Specifically, for the regions Holland-Midden, Twente and Groningen, the mothership protocol performs worse in 68.4% of the instances. Over all instances, the fraction of IAT patients that need to be transferred from PSC to CSC in the optimal model is 59.4% on average with a standard deviation of 25.1%. This shows that the required number of transfers will be considerable, but strongly depends on the region and health-related parameters.

Sensitivity of parameters So far, we primarily focused on the impact of the regions on the performance. Below, we further explore the impact of $P(IAT)$ and the total demand P ; the impact of the minimum IAT requirement r_{IAT} is related to the impact of $P(IAT)$ and P . In Figs. 8 and 9 you may find the mean (line) and interquartile range (shaded area) of the total SDST for the two practical protocols relative to the optimal protocol against $P(IAT)$ and P , respectively. As may be expected, the performance of the drip-and-ship protocol decreases as $P(IAT)$ increases, whereas the performance of mothership improves. In fact, for $P(IAT) = 60\%$ we see that SDST for mothership is close to optimal. For small values of $P(IAT)$, e.g. when $P(IAT) = 20\%$, drip-and-ship typically performs better than mothership, but the performance is not necessarily close to optimal. Specifically, in the case with small patient

Table 5 Fraction of instances (in %) in which mothership outperforms drip-and-ship per region for different values of $P(IAT)$

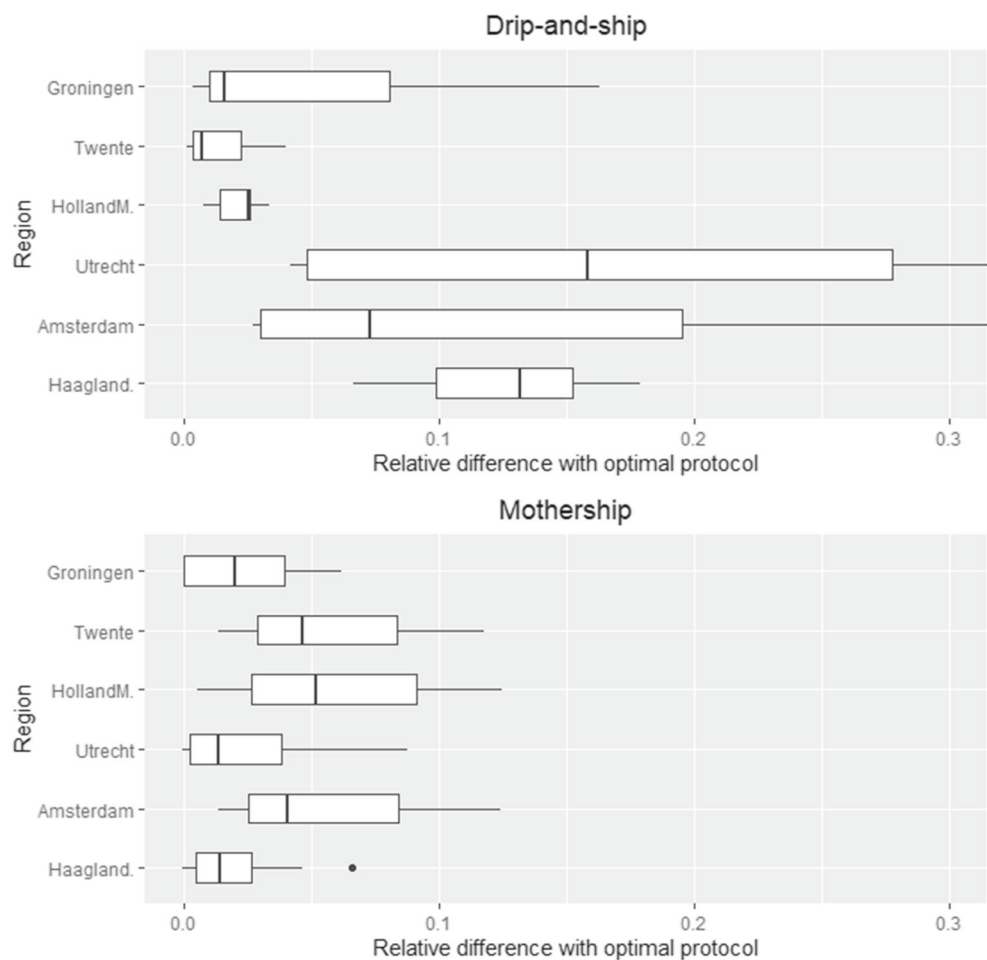
Region	Rurality	P(IAT)					Total
		20%	30%	40%	50%	60%	
Haaglanden	urban	100.0	100.0	100.0	100.0	100.0	100.0
Amsterdam	urban	16.7	14.3	62.5	66.7	100.0	56.4
Utrecht	mixed	50.0	57.1	100.0	100.0	100.0	84.6
Holland-Midden	mixed	0.0	0.0	0.0	0.0	100.0	23.1
Twente	rural	0.0	0.0	0.0	33.3	33.3	15.4
Groningen	rural	16.7	14.3	25.0	100.0	100.0	56.4

volumes ($P = 300$), the SDST of drip-and-ship for Utrecht is still 41.9% worse than optimal.

Figure 9 shows the performance of the protocols relative to optimal as the total number of patients P varies. For $P = 300$, the performance of drip-and-ship varies considerably. Specifically, the performance for the regions Utrecht and Amsterdam is quite poor in that case, and are at least 37% (Utrecht) and 31% (Amsterdam) worse than optimal for all values of $P(IAT)$. This can be explained by the in-hospital delay that becomes rather big when there are quite some stroke

centres (6 and 5 for Utrecht and Amsterdam, respectively) and the total patient volume is small. These examples show that drip-and-ship is vulnerable when the number of patients per stroke centre becomes small; the order of magnitude also clearly depends on the function chosen for b_{IVT} when patient volumes become small. For mothership, the relative performance compared to optimal decreases when P increases. Again, this follows from in-hospital delay; as P is larger, the optimal protocol will use more PSCs for its allocation without excessive in-hospital delays for IVT.

Fig. 7 Relative differences with optimal protocol for drip-and-ship (top) and mothership (bottom)






Fig. 8 Mean relative difference with the optimal protocol against $P(IAT)$ (the area between the 25% and 75% percentile is shaded)

Remark 5 Although the drip-and-ship and mothership protocols may show close to optimal performance in some instances, it seems difficult to give any reasonable performance guarantees. For instance, consider a simplified situation with $P(IAT) = 0$, such that we may focus on IVT. Assume that there are 3 demand locations ($I = 3$) and two possible PSC locations ($J = 2$) that are the same as demand locations 1 and 2. Moreover, let $d_{12} = d_{21} = D$, $d_{31} = D - \delta$ and $d_{32} = \delta$ for some $\delta > 0$ sufficiently small and D large. Consider the following demands: $w_1 = r_{IVT} - \epsilon$, $w_2 = r_{IVT}$, and $w_3 = \epsilon$ for some small $\epsilon > 0$.

In the optimal allocation, both PSC locations are open and patients of locations 1 and 3 are assigned to PSC location 1 such that the minimum volumes of IVT are met for both PSCs. The optimal total travel time is then $\epsilon(D - \delta)$, which is only due to patients from location 3. When applying a drip-and-ship protocol it is not possible to open both locations, as PSC location 2 would be the nearest PSC to demand location 3. Hence, for drip-and-ship it is only feasible to open one location without violating the minimum volume requirements. Opening location 2 as a PSC, then provides the smallest total travel time of $(r_{IVT} - \epsilon) \times D + \epsilon \times \delta = r_{IVT}D - \epsilon(D - \delta)$. Observe that both the absolute and

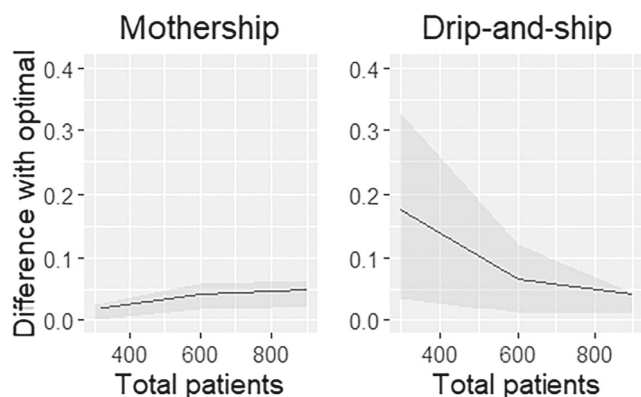


Fig. 9 Mean relative difference with the optimal protocol against P (the area between the 25% and 75% percentile is shaded)

relative difference in total travel time between optimal and drip-and-ship explode when $D \rightarrow \infty$ and $\epsilon \downarrow 0$.

5.4 Computational time

The numerical experiment is run on an Intel Core i7-4770k CPU @ 3.50GHz with 24 GB RAM. In this experiment, a single instance of the model takes approximately one minutes; with 270 instances and three models (mothership, drip-and-ship, and optimal) the total experiment takes roughly 13 hours to complete.

6 Conclusion and discussion

In this paper, we model the joint decision of PSC and CSC locations and the allocation of acute stroke patients in a region as a two-level hierarchical facility location problem. Specifically, we include the impact of the volume of IVT patients on in-hospital delays in our MILP, and compare the optimal model with the protocols in practice: mothership and drip-and-ship.

From our numerical experiments we see that the performance of the protocols depends on a variety of elements. In general, mothership performs better than drip-and-ship in urban areas, in particular in case of more centrally located CSCs; vice versa, drip-and-ship performs well in rural areas and stroke centres spread over the region. Additionally, for $P(IAT) \leq 30\%$ drip-and-ship seems preferable, unless the total patient volume per PSC becomes rather small, whereas mothership performs excellent for $P(IAT) \geq 50\%$. On average, mothership and drip-and-ship are 3.9% and 8.6% worse than optimal, whereas drip-and-ship seems more sensitive to parameter choices.

In our experiments, the optimal protocol leads to a decrease between 0 and 41.9 percentage points in SDST for a region. This corresponds to a gain between 0 and 18.9 minutes per treated patient, based on an average SDST of 45 minutes. A measure to evaluate the impact of medical treatments in economical terms is QALY's (quality-adjusted life years) [33]. Within the first six hours of onset, every hour of delay results in an average loss of 0.77 QALY's for IAT patients [36]. Assuming linearity, this results in an improvement between 0 and 0.24 QALY's for an individual IAT patient. For a region with 150 IAT patients per year, with a value of € 50,000 per QALY [35], the value of medical intervention in acute stroke care has the potential to gain a value up to € 1,800,000 per year using the optimal model, although the potential strongly depends on the region and the parameters.

The model precisely prescribes which stroke centres should become PSCs and CSCs and how patients should be allocated depending on the location of onset. As a consequence,

the flow of stroke patients, including patient transfers, will typically change and stroke centres will be faced with different patient volumes. As patient volume is key to any hospital, the potentially required modifications will be a subject of debate. Due to the rather recent introduction of IAT, there is now a practical need to re-establish the allocation of ischemic stroke patients. The insights from the numerical experiments provide an initial estimate of a proper organization and the potential gain. Moreover, the application of the optimal model for a specific region and situation should facilitate the debate about locations of PSCs and CSCs and the corresponding allocation. Even if the optimal model is not fully implemented, it shows the direction to look for when reorganizing acute stroke care.

Finally, there are some limitations to the model that may be studied in the future. The regional layout and the parameters, such as $P(IAT)$ and b_{IVT} , may have a considerable impact on which protocol is better, but these parameters are difficult to estimate in current practice. Also, the value of $P(IAT)$ may change in the future and may depend on the population demographics and IAT eligibility criteria.

The overall observations from our numerical experiments provide insight in the differences between the two practical protocols, but for any region it would be preferred to apply an optimization for finding the actual optimal allocation.

Also, in the current model we minimize the total time SDST, but the actual criterion is the condition of the patient after IVT and/or IAT. From a computational perspective, it may be observed that the problem is NP-hard. This implies that the size of the region that can be calculated to optimality using an MILP formulation is limited.

Acknowledgments This work was supported by The Netherlands Brain Foundation (project number HA2015.01.02), The Dutch, Health Care Insurers Innovation Foundation (project number 3240) and Health Holland (project number: LSHM16041-HOOB). Moreover, we would like to thank Geert-Jan Kommer from the Dutch National Institute for Public Health and the Environment (RIVM) for providing the travel time and demand data.

Compliance with Ethical Standards

Conflict of interests The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article. In accordance with the policies of the institutions of the LUMC and VU Amsterdam, no ethics approval was needed for the used data.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds

the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A : Mothership model

In this section, we provide the MILP formulation for the mothership protocol, i.e., the locations of CSC facilities. Since the allocation between stroke centres is no longer possible, Constraints (4) should be removed and the problem is now a single-level system. Due to the absence of PSCs, the notations may be simplified. However, we choose to use the same notation for consistency with the models for the other protocols.

Observe that Constraints (25) makes sure that each location either provides both IVT and IAT, or no treatment of acute stroke. Constraints (30) are similar to the drip-and-ship protocol and guarantees that each patient is allocated to the nearest location.

$$\text{Minimize } Z_{moth} = \sum_{i=1}^{|I|} \sum_{j=1}^{|J|} z_{ij} + y_{ij} w_i (d_{ij} + b_{j,IAT} P(IAT)) \tag{20}$$

$$\text{s.t. } \sum_{j=1}^{|J|} y_{ij} = 1, \forall i \tag{21}$$

$$\sum_{i=1}^{|I|} y_{ij} w_i \geq r_{IVT} h_{j,IVT} \forall j \tag{22}$$

$$\sum_{i=1}^{|I|} y_{ij} w_i P(IAT) \geq r_{IAT} h_{j,IAT} \forall j \tag{23}$$

$$\sum_{i=1}^{|I|} y_{ij} \leq M h_{j,IVT} \forall j \tag{24}$$

$$h_{j,IVT} = h_{j,IAT} \forall j \tag{25}$$

$$z_{ij} \leq M y_{ij} \forall i, j \tag{26}$$

$$z_{ij} \leq b_{j,IVT} w_i \forall i, j \tag{27}$$

$$z_{ij} \geq b_{j,IVT} w_i - (1 - y_{ij}) M \forall i, j \tag{28}$$

$$b_{j,IVT} \geq \sum_{i=1}^{|I|} \alpha_n w_i y_{ij} + \beta_n \forall n, i, j \tag{29}$$

$$\sum_{k \in J_j^*} h_{k,IVT} \leq (1 - y_{ij}) M \forall i, j \tag{30}$$

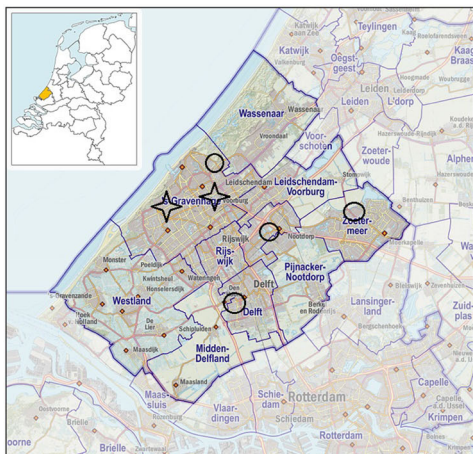
$$y_{ij} \in \{0, 1\} \forall i, j \tag{31}$$

$$h_{j,c} \in \{0, 1\} \forall j, c \tag{32}$$

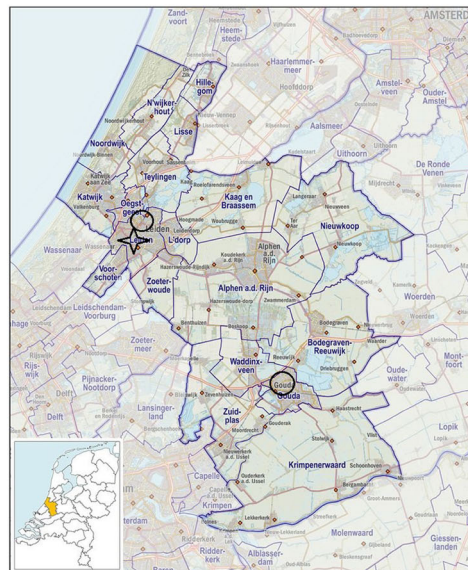
$$z_{ij} \geq 0 \forall i, j \tag{33}$$

$$b_{j,IVT} \geq 0 \forall j \tag{34}$$

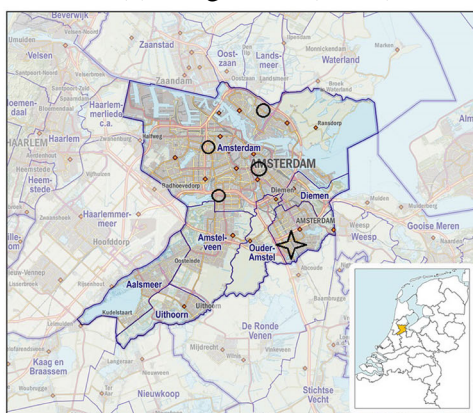
Appendix B : Region structures



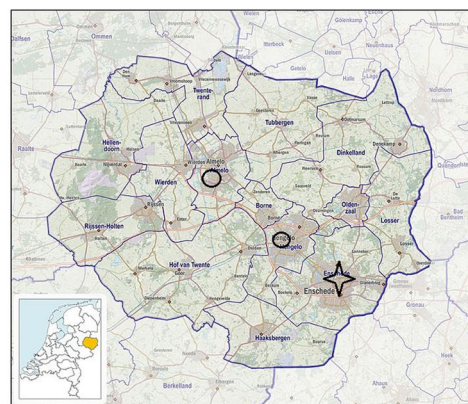
(a) Haaglanden (urban)



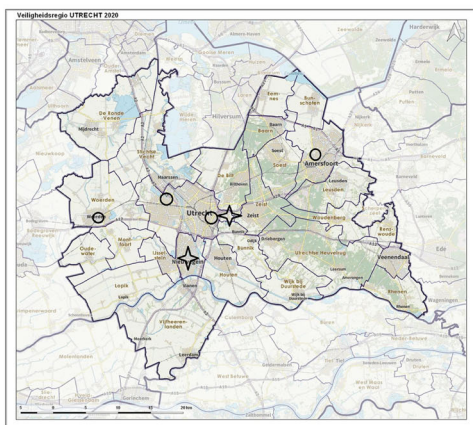
(d) Holland-Midden (mixed)



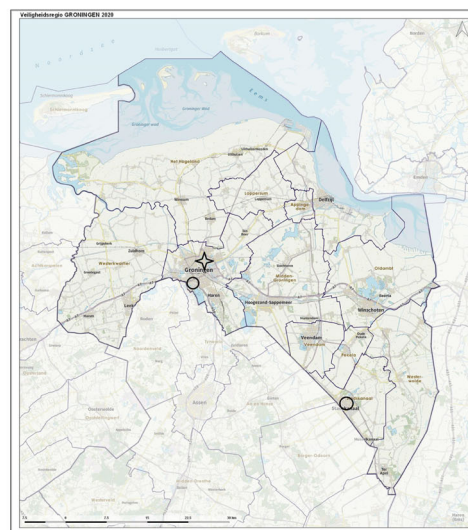
(b) Amsterdam-Amstelland (urban)



(e) Twente (rural)



(c) Utrecht (mixed)



(f) Groningen (rural)

Fig. 10 Lay-out of the six regions

References

1. Goyal M, Menon B, van Zwam WH et al (2016) Endovascular thrombectomy after large-vessel ischaemic stroke: a meta-analysis of individual patient data from five randomised trials. *Lancet* 387:1723–1731
2. Berkhemer OA, Fransen PS, Beumer D et al (2015) A randomized trial of intraarterial treatment for acute ischemic stroke. *N Engl J Med* 372:11–20
3. Goyal M, Demchuk AM, Menon B et al (2015) Randomized assessment of rapid endovascular treatment of ischemic stroke. *N Engl J Med* 372:1019–1030
4. Saver JL, Goyal M, Bonafe A et al (2015) Stent-retriever thrombectomy after intravenous t-PA vs. t-PA alone in stroke. *N Engl J Med* 372:2285–2295
5. Campbell BC, Mitchell PJ, Kleinig TJ et al (2015) Endovascular therapy for ischemic stroke with perfusion-imaging selection. *N Engl J Med* 372:1009–1018
6. Jovin TG, Chamorro A, Cobo E et al (2015) Thrombectomy within 8 hours after symptom onset in ischemic stroke. *N Engl J Med* 372:2296–2306
7. Badhiwala JH, Nassiri F, Alhazzani W et al (2015) Endovascular thrombectomy for acute ischemic stroke: a meta-analysis. *JAMA* 314:1832–1843
8. Nicholl J, West J, Goodacre S, Turner J (2007) The relationship between distance to hospital and patient mortality in emergencies: an observational study. *Emerg Med J* 24:665–668
9. Koster GT, Nguyen TTM, Groot AE et al (2018) A reduction in time with electronic monitoring in stroke (ARTEMIS): study protocol for a randomised multicentre trial. *BMJ Open* 8:e020844
10. Ortiz-Astorquiza C, Contreras I, Laporte G (2017) Multi-level facility location problems. *Eur J Oper Res* 267:791–805
11. Sahin G, Sural H (2007) A review of hierarchical facility location models. *Comput. Oper. Res.* 34:2310–2331
12. Ahmadi-Javid A, Seyedi P, Syam SS (2017) A survey of healthcare facility location. *Computers & Operations Research* 79:223–263
13. Van den Berg PL, van Essen JT, Harderwijk EJ (2016) Comparison of static ambulance location models. In: 3rd IEEE International Conference on Logistics Operations Management, pp 1–10
14. Groot AE, van Schaik IN, Visser MC et al (2016) Association between iv thrombolysis volume and door-to-needle times in acute ischemic stroke. *J Neurol* 263:807–813
15. Parikh NS, Chatterjee A, Diaz I. et al (2018) Modeling the impact of interhospital transfer network design on stroke outcomes in a large city. *Stroke* 49:370–376
16. Lima FO, Silva GS, Furie KL et al (2016) Field assessment stroke triage for emergency destination: a simple and accurate prehospital scale to detect large vessel occlusion strokes. *Stroke* 47:1997–2002
17. Milne MS, Holodinsky JK, Hill MG et al (2017) Drip ‘n ship versus mothership for endovascular treatment: modeling the best transportation options for optimal outcomes. *Stroke* 48:791–794
18. Holodinsky JK, Patel AB, Thornton J et al (2018) Drip and ship versus direct to endovascular thrombectomy: The impact of treatment times on transport decision-making. *European Stroke Journal* 3:126–135
19. Kuhrij LS, Wouters MWJM, van den Berg-Vos RM et al (2018) The Dutch Acute Stroke audit: Benchmarking acute stroke care in the Netherlands. *European Stroke Journal* 3:361–368
20. Koster GT, Nguyen TTM, van Zwet EW et al (2019) Clinical prediction of thrombectomy eligibility: A systematic review and 4-item decision tree. *International Journal of Stroke* 14:530–539
21. Shahid R, Bertazzon S, Knudtson ML, Ghali WA (2009) Comparison of distance measures in spatial analytical modeling for health service planning. *BMC Health Services Research* 9:200
22. Tawil SE, Cheripelli B, Huang X et al (2016) How many stroke patients might be eligible for mechanical thrombectomy? *European Stroke Journal* 1:264–271
23. Mathai AM, Moschopoulos P, Pederzoli G (1999) Random points associated with rectangles. *Rendiconti del Circolo Matematico di Palermo* 48:163–190
24. Moltchanov D (2012) Distance distributions in random networks. *Ad Hoc Netw* 10:1146–1166
25. Fransen PSS, Beumer D, Berkhemer OA et al (2014) MR CLEAN, a multicenter randomized clinical trial of endovascular treatment for acute ischemic stroke in the Netherlands: study protocol for a randomized controlled trial. *Trials* 15:343
26. English JD, Yavagal DR, Gupta R et al (2015) Mechanical thrombectomy-ready comprehensive stroke center requirements and endovascular stroke systems of care: recommendations from the Endovascular Stroke Standards Committee of the Society of Vascular and Interventional Neurology (SVIN). *Interventional Neurology* 4:138–150
27. Stone RE (1991) Some average distance results. *Transp Sci* 25:83–90
28. Ismail M, Armoiry X, Tau N et al (2019) Mothership versus drip and ship for thrombectomy in patients who had an acute stroke: a systematic review and meta-analysis. *Journal of Neurointerventional Surgery* 11:14–19
29. Bray BD, Campbell J, Cloud GC et al (2013) Bigger, faster? Associations between hospital thrombolysis volume and speed of thrombolysis administration in acute ischemic stroke. *Stroke* 44:3129–3135
30. Hubert GJ, Meretoja A, Audebert HJ et al (2016) Stroke thrombolysis in a centralized and a decentralized system (Helsinki and Telemedical Project for Integrative Stroke Care Network). *Stroke* 47:2999–3004
31. Strbian D, Ahmed N, Wahlgren N et al (2015) Trends in door-to-thrombolysis time in the safe implementation of stroke thrombolysis registry: effect of center volume and duration of registry membership. *Stroke* 46:1275–1280
32. White PM, Bhalla A, Dinsmore J et al (2017) Standards for providing safe acute ischaemic stroke thrombectomy services. *Clinical radiology* 72:175–e1
33. Weinstein MC, Torrance G, McGuire A (2009) QALYs: The basics. *Value in Health* 12:S5–S9
34. Intrevador P, Verter V, Tremblay L (2019) Patient-centric design of long-term care networks. *Health Care Management Science* 22:376–390
35. Pomp M (2010) Een beter Nederland: de gouden eieren van de gezondheidszorg. Amsterdam: Balans
36. Kunz WG, Almekhlafi M, Menon B et al (2018) Lifetime quality of life and cost consequences of treatment delays in endovascular thrombectomy for stroke based on hermes data. *Journal of NeuroInterventional Surgery* 10:A1–A2
37. Venema E, Groot AE, Lingsma HF et al (2019) Effect of interhospital transfer on endovascular treatment for acute ischemic stroke. *Stroke* 50:923–930
38. Van den Berg PL, van Essen JT (2019) Comparison of static ambulance location models. *International Journal of Logistics Systems and Management* 32:292–321
39. Güneş ED, Melo T, Nickel S (2019) Location problems in healthcare. In: *Location science*, 657–686, Springer, Cham
40. Laporte G, Nickel S, da Gama FS (2019) *Location science*, vol. 528, 2nd edition, Springer, Berlin
41. Smith HK, Harper PR, Potts CN (2013) Bicriteria efficiency/equity hierarchical location models for public service application. *Journal of the Operational Research Society* 64:500–512
42. Verter V, Lapiere SD (2002) Location of preventive health care facilities. *Ann Oper Res* 110:123–132

43. Chao X, Liu L, Zheng S (2003) Resource allocation in multisite service systems with intersite customer flows. *Manag Sci* 49:1739–1752
44. Zhang Y, Berman O, Verter V (2009) Incorporating congestion in preventive healthcare facility network design. *Eur J Oper Res* 198:922–935
45. Zhang Y, Berman O, Marcotte P, Verter V (2010) A bilevel model for preventive healthcare facility network design with congestion. *IIE Trans* 42:865–880
46. Vidyarthi N, Kuzgunkaya O (2015) The impact of directed choice on the design of preventive healthcare facility network under congestion. *Health care management science* 18:459–474
47. Ahmadi-Javid A, Ramshe N (2020) A stochastic location model for designing primary healthcare networks integrated with workforce cross-training. *operations research for health care* 24
48. Narula SC, Ogbu UI (1979) An hierarchal location-allocation problem. *Omega* 7:137–143
49. Mestre AM, Oliveira MD, Barbosa-Póvoa AP (2015) Location–allocation approaches for hospital network planning under uncertainty. *European Journal of Operational Research* 240:791–806
50. Cardoso T, Oliveira MD, Barbosa-Póvoa A, Nickel S (2015) Introducing health gains in location-allocation models: A stochastic model for planning the delivery of long-term care. In: *Journal of Physics: Conference Series* 616, 012007
51. Cardoso T, Oliveira MD, Barbosa-Póvoa A, Nickel S (2016) Moving towards an equitable long-term care network: a multi-objective and multi-period planning approach. *Omega* 58:69–85

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.