A Genetic Algorithm-based strategic planning framework for optimising accessibility and costs of General Practices in Northland, New Zealand

Fulvio D. Lopane[a] and Melanie Reuter-Oppermann[b] and Andrea Raith[c] and Daniel J Exeter[c] and Ilze Ziedins[c] and Richard Dawson[d]

[a]Centre for Advanced Spatial Analysis, University College London, United Kingdom; [b]Technical University of Darmstadt, Germany; [c]University of Auckland, New Zealand; [d]School of Engineering, Newcastle University, Newcastle upon Tyne, United Kingdom

ARTICLE HISTORY
Compiled January 16, 2023

ABSTRACT
Shortage of general practitioners (GP) is a challenge worldwide, not only in Europe, but also in countries like New Zealand. Providing primary care in rural areas is especially challenging. In order to support decision makers, it is necessary to first assess the current GP coverage and then to determine different scenarios and plans for the future. In this paper, we first present a thorough overview of related literature on locating GP practices. Second, we propose an approach for assessing the GP coverage and determining future GP locations based on a genetic algorithm framework. As a use case we have chosen the rural New Zealand region of Northland. We also perform a sensitivity analysis for the main input parameters.

KEYWORDS
General practitioners; Spatial optimisation; New Zealand; Decision support

1. Introduction

A majority of countries worldwide suffer from a shortage of general practitioners (GPs). Providing primary care in rural areas is especially challenging, while the definition for rural is significantly different between the countries. In Germany, for example, a region is defined as rural if 140 inhabitants are living in one km^2. In New Zealand, the average population density is only 18 inhabitants per km^2 (Worldometer (2021)). Even though the healthcare systems in the two countries differ significantly, primary care practices need a sufficient number of patients to be profitable and attractive for GPs to work in. At the same time, primary care is crucial for people's welfare. Patients benefit from having easy and timely access to GPs in close proximity, often to their home, but sometimes to their place of work. Especially in rural areas this is difficult to ensure as GPs also need sufficient patients in their panel, i.e. the number of patients per GP, to cover their costs. Often, too few GPs are available to fill open spots, and rural practices can be even less attractive. Due to the shortage of GPs in New Zealand (Goodyear-Smith and Janes (2008)) there is a need for a wise use of resources, for example by

CONTACT M. Reuter-Oppermann. Hochschulstr. 1, 64289 Darmstadt, Germany. Email: oppermann@is.tu-darmstadt.de

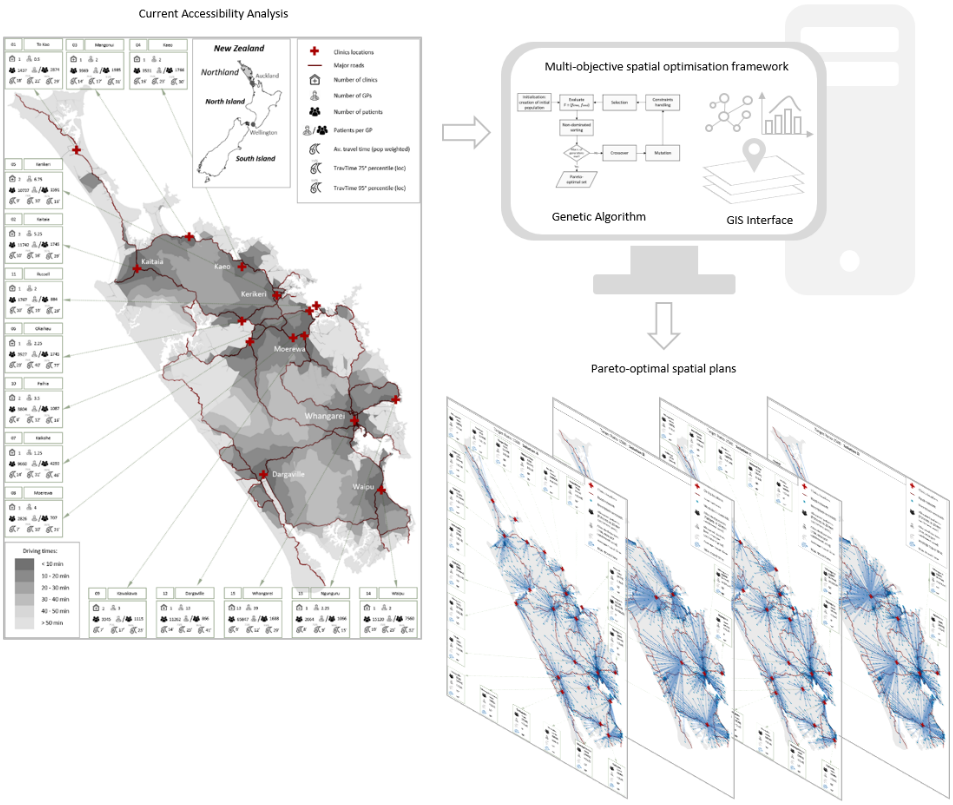

Figure 1. Summary of the contribution

reasonably locating GPs and practices in order to minimise traveling distances for patients and assuring a “sufficient” number of patients per GP. In order to support decision makers with location planning, it is necessary to first assess the current GP coverage.

This paper assesses as an example the coverage of GPs for the New Zealand region of Northland. Northland is a comparably rural area in the very north of New Zealand’s North Island (Zhao, Ameratunga, Lee, Browne, and Exeter (2019)). Its approximately 150.000 inhabitants live across 12.500 km^2, leading to an average of 12 inhabitants per km^2. By adapting a strategic planning and spatial accessibility framework (Lopane, Barr, James, and Dawson (2019)), we present a decision support system to determine potential GP locations together with the resulting coverage areas and expected driving times of patients, assuming that patients attend the closest GP in relation to their home. Besides the travel costs, a genetic algorithm also minimises the infrastructure costs, thereby determining Pareto-optimal spatial plans. By varying the target ratio (TR) for panel sizes and the number of full-time doctors per practice a sensitivity analysis is performed and the usability of the framework to support decision making about practice locations is demonstrated. The contribution is summarised in Figure 1.

The remainder of the paper is structured as follows. Section 2 reviews the background and context of the problem. Then, in Section 3, the methodology is presented. First, the status quo of primary care in the rural area of Northland, New Zealand, is assessed. Then, using a spatial optimisation framework, optimal location strategies for primary care practices are identified under different scenarios. The findings are discussed in Section 4 and the paper finishes with a conclusion and an outlook on future research in Section 5.

2. Foundations

In this section, existing publications on siting GP practices from the area of operations research as well as relevant papers from health services research and urban planning are summarised. While this is not a review paper, we still aim at giving a good overview, as to the best of our knowledge a proper review paper on locating primary care practices is still missing in the literature. In addition, a short description of the New Zealand healthcare system is given to provide necessary background information for the case study presented in this work.

2.1. Literature review

While facility location planning is a major research area in operations research, literature on locating primary care practices is still rather limited. Siting other types of healthcare facilities has been studied more extensively. Overviews on healthcare facility location planning can be found in Güneş, Melo, and Nickel (2019), Daskin and Dean (2005) and Farahani, SteadieSeifi, and Asgari (2010). Güneş et al. (2019) for example define three objectives for locating healthcare facilities: (1) minimise the access cost for patients (e.g. travel cost, distance or travel time), (2) maximise covered demand, and (3) maximise equity in access. Rahman and Smith (2000) investigate the use of location-allocation models in health service development planning in developing countries in their review .

Abernathy and Hershey (1972) addressed the problem of locating primary care facilities. They discussed and applied different objective functions, e.g. maximising the utilisation or minimising the distances. Parker and Srinivasan (1976) studied the location of additional primary care facilities to expand an existing care network in a rural region and proposed a five-step approach to address it. They use a heuristic to determine those facilities that maximise the total incremental benefit subject to predefined cost constraints. Hillsman (1980) reported the outcome of a study to locate primary care facilities in rural Iowa, USA. He described the same challenge for Iowa that we see in many countries worldwide today: retiring GPs cannot be replaced at the same rate leading to a decreasing number of GPs and potentially longer distances for patients, especially in rural areas. He used a (simple) bi-objective location-allocation model as a variation of the p-median model that minimises the travelling distances to the facilities and maximises the utilisation of existing resources. Due to the size of the problem, he developed an algorithm to determine 30 solutions with different trade-offs for the two objectives. The primary care network development in Ontario, Canada, from a regulator's point of view was studied by Graber-Naidich, Carter, and Verter (2015). They propose a mixed integer linear programming model that can be used to analyse the outcomes of locating different types of primary care facilities on the overall cost, accessibility and appropriateness of care provided. The model uses a deprivation index that correlates with increased health needs and barriers to care. Reuter-Oppermann, Nickel, and Steinhäuser (2019) applied three operations research models for locating GP practices in a German district. Models were developed under two main basic requirements: (1) one practice is to be located that can be reached by as many inhabitants as possible, and (2) locate practices to cut down the driving time for every district's inhabitants to the next practice location to less than 15 minutes. Input data included the demand (population), driving times and the current GP locations. Güneş, Yaman, Çekyay, and Verter (2014) introduced an integer programming model for planning primary care facility networks in Turkey that integrates the planner's perspective with patient and physician preferences. Due to the potentially contradictory objectives, the authors investigate trade-offs and effects of parameter changes. Panagiotis Mitropoulosa (2013) developed a planning process for restructuring the network of primary care facilities by repeatedly employing data envelopment analysis (DEA) and solving a multi-objective location-allocation model. While DEA measures the technical efficiency of existing facilities, the location model determines the appropriate number and locations of facilities in the network. The model includes the decision about potential upgrades or closings of facilities. The authors applied their approach to the region of the Peloponnese in South Greece.

Ahmadi-Javid and Ramshe (2020) used stochastic programming for locating and staffing health posts as part of the primary care network that offer a limited number of specialised care types. They applied their model to the city of Tehran. Tien and El-Tell (1984) focused on primary care in developing countries and present a hierarchical location-allocation model for primary care facilities in Jordan. The ILP model locates two facility types minimising the travelling distance for patients while trying to cover as much demand as possible by the facilities with the lower care level. Especially in developing countries with a significant GP shortage, potentially low population density, but long distances between villages, mobile primary care practices can be a good solution to provide primary care to the whole country. Then, the location problem is combined with a routing problem, as for example proposed by Hodgson, Laporte, and Semet (1998).

A related problem is the districting of healthcare regions. In their review on districting problems in healthcare, Yanık and Bozkaya (2020) also presented districting models for primary care. The idea is that all patients within one district are served together. If primary care is centralised and districts are small enough, patients are served by one practice per region and the districting problem is equivalent to a simple location-allocation problem. A multi-objective districting problem for partitioning the healthcare system of Paraná State in Brazil has been developed by Steiner, Datta, Neto, Scarpin, and Figueira (2015). They considered the three objectives (1) maximising the population homogeneity in the districts, (2) maximising the variety of medical procedures offered in the districts and (3) minimising the distances to be travelled by patients. The authors propose a genetic algorithm for solving the problem. Shortt, Moore, Coombes, and Wymer (2005) presented an approach to determine catchment areas for GP practices, i.e. locations of patients attending certain GP practices, taking patient flow into account.

Preventive care facilities usually offering a limited number of dedicated care services are closely related to primary care. Verter and Lapierre (2002), for example, presented an analytical framework based on a mathematical formulation for the preventive health care facility location problem for determining the optimal design of a preventive health care facility network with the aim to maximise the patients' participation in prevention programs. Y. Zhang, Berman, and Verter (2009) incorporated the issue of congestion in preventive care facilities into the model as a queuing model. As the resulting optimisation model is non-linear, the authors present heuristics for solving the problem. Y. Zhang, Berman, Marcotte, and Verter (2010) presented a bilevel non-linear optimisation model for the same problem that is partly solved as a convex optimisation problem and with the help of a Tabu Search heuristic. In addition,

Y. Zhang, Berman, and Verter (2012) studied the impact of client choice behaviour on the design of a preventive care facility network and the resulting level of patient participation in preventive care programs.

Community Health Centers (CHC) in the United States offer health care services for people living in medically underserved communities, including primary and preventive healthcare, outreach and dental care, as well as several welfare services. Griffin, Scherrer, and Swann (2008) presented an optimisation model that determines both locations and service offerings for CHCs taking demand, costs and facility sizes into account.

Often, healthcare networks and facilities already exist. The aim is then to use location planning to restructure and improve the network, for example by locating new facilities or changing locations. This is not always easy to implement, but takes a long time and involves high costs. As a first step to discover the actual problems and the current status of the healthcare network, researchers have proposed the use of spatial analysis of healthcare facilities.

Baum, Bergwall, and Reeves (1975) presented the "HEAlth-care Delivery Simulator for Urban Population" (HEADSUP II) to analyse primary health care delivery in an urban or suburban setting. Following a similar idea, Standridge, Alan, Pritsker, and Delcher (1978) also developed a simulation in 1978 to study the primary healthcare system of Indiana, USA, through the year 2000. In order to measure spatial accessibility in primary care, Luo and Qi (2009) proposed an enhancement of the two-step floating catchment area (2SFCA) method, addressing the problem of uniform access within the catchment areas and accounting for distance decay by applying weights to different travel time zones. Schuurman, Berube, and Crooks (2010) used a modified gravity model to measure spatial access to primary care practices in the Canadian province of Nova Scotia. The model incorporated a distance decay function and allowed for people to travel across artificial census boundaries to access a primary care practice.

In their research, several authors have addressed access to primary care in Australia, especially spatial access disparities in rural areas, and studied different measures and catchment geographies (Butler, Petterson, Bazemore, and Douglas (2010); Duckett, Breadon, and Ginnivan (2013); Mazumdar et al. (2019, 2016); Mazumdar, Feng, Konings, McRae, and Girosi (2014); McGrail and Humphreys (2009, 2015)).

A review on concepts, methods and challenges for spatial accessibility of primary care was published by Guagliardo (2004).

The accessibility of public hospitals in New Zealand was studied by Brabyn and Skelly (n.d., 2002). They found that the average travel time to a hospital is around 20 minutes, to the closest tertiary hospital even 90 minutes.

For analysing and improving locations of care facilities, using a geographical information system (GIS) is advantageous. Several authors have studied the use of GIS for locating new or analysing the access to existing GP practices (Bullen, Moon, and Jones (1996); Crooks and Schuurman (2012); Foley and Darby (2002); Luo and Wang (2003)) or other healthcare / public facilities (Foley (2002); Gordon and Womersley (1997); Lapierre, Myrick, and Russell (1999); McLafferty (2003); Murad (2004); Rosero-Bixby (2004); Walsh, Page, and Gesler (1997)). A review on the interaction between GIS and location science has been published by Church (2002), for example.

GIS-based decision support systems seem especially promising for the use in practice. Ribeiro and Antunes (2002) as well as Reuter-Oppermann, Rockemann, and Steinhäuser (2017) proposed to combine a GIS with a mathematical solver to determine the optimal solution to a location problem and present the results with the help of GIS, while Ribeiro and Antunes studied public facilities in general and Reuter-Oppermann et al. focus on GP practices.

Due to the importance of analysing existing and determining new and improved GP locations, this paper combines and extends the findings of previous publications by presenting a spatial framework that allows the integration of planning approaches and displays locations and catchment areas together with relevant information on distances and expected number of patients per GP in a GIS.

2.2. New Zealand healthcare system

New Zealand is divided into 20 District Health Boards (DHBs) that are responsible for providing or funding the provision of health services in their district. They have service agreements with private and NGO providers, e.g. GPs, to provide care for the New Zealand population who in turn might need to pay / co-pay for these services (Ministry of Health NZ (2021a)).

Large parts of the New Zealand healthcare system, especially hospital care, are funded by general taxation. Therefore, these services are free or subsidised for people who are eligible for publicly funded healthcare in New Zealand. This also includes emergency departments (EDs) at public hospitals that can be accessed freely at no cost. By contrast, patients are required to pay for their visits, and their medicines/drugs to treat the illness(es) carry a co-payment, as primary care is not fully funded by the government (Southern Health (2021)).

In New Zealand, there are around 40 public hospitals. GPs' practices are usually open during business hours from Monday to Friday. Outside these hours, practices should make arrangements for their patients to still receive care (Ministry of Health NZ (2021a)).

Patients in New Zealand can choose the general practice that they want to visit. On the other hand, GPs can decide whether they can add another patient to their panel, or reject new patients if they already have taken on too many. Often, the chance is higher for a patient to be accepted if the practice is located in the neighbourhood where they live. If a patient cannot be taken on by a GP, the practice refers the patient to their primary health organisation (PHO) for help with finding another practice. The PHO may put a patient on a waiting list and arrange for them to get care in the meantime. The New Zealand Medical Council has a register of practising GPs and advice on choosing a GP (Ministry of Health NZ (2021a)).

Once a general practice is found, it is also possible to enrol with this practice free of charge. Then, the care will be subsidised by the Government and patients only pay a reduced consultation fee. GPs normally charge a higher fee, often called a casual rate, for patients that are not enrolled with their practice. General practices can only enrol people who are eligible for publicly funded primary health services, though. General practices are private businesses that can set their own fees for consultations and other health services. While the fees charged must be within a certain threshold agreed to by the DHBs and PHOs, practices determine the level of co-payment. The list of general practices and their fees can be found on the websites of the local DHBs and/or of the PHOs. All children under 13 are eligible for free general practice visits, both during the day and after-hours, but not all GPs offer these free visits. The PHOs runs the Very Low Cost Access (VLCA) programme that GPs can join to receive extra funding from the Government to keep their fees at low levels for all enrolled patients. Therefore, there are significant differences in fees between practices (Ministry of Health

Table 1. Northland current GP practices accessibility analysis (clinics grouped by location). Travel times (TT) are calculated basing of road network edges' allowed speeds.

ID	Location	N of clinics	FTE GPs	Popula-tion	Patients / GPs	Average TT [min]	75% quantile TT [min]	95% quantile TT [min]
01	Te Kao	1	0.5	1437	2874	17	21	29
02	Kaitaia	2	5.25	11742	2237	13	16	39
03	Mangonui	1	2	3969	1985	11	17	31
04	Kaeo	1	2	3531	1766	16	23	30
05	Kerikeri	2	6.75	10737	1591	9	10	16
06	Okaihau	1	2.25	3927	1745	28	40	77
07	Kaikohe	1	2.25	9660	4293	18	31	46
08	Moerewa	1	4	2826	707	7	10	21
09	Kawakawa	2	3	3345	1115	10	17	23
10	Paihia	2	3.5	3804	1087	8	12	16
11	Russell	1	2	1767	884	11	19	29
12	Dargaville	1	13	11262	866	17	25	41
13	Ngunguru	1	2.5	2664	1066	8	9	15
14	Waipu	1	2	15120	7560	19	25	32
15	Whangarei	13	39	65847	1688	9	12	29

NZ (2021b)).

The GP keeps the enrolled patients' medical records, but any health professional involved in a patient's care can look at the record. Also, patients can get access to their record at any time. When changing the GP and enrolling with a new practice, records can be transferred from the old practice to the new one (Ministry of Health NZ (2021b)).

3. Methodology

The ultimate goal of the study is to demonstrate how spatial optimisation can be used by local authorities and healthcare providers to make more effective use of healthcare resources. Our study consists of a 2-step approach that first involves an assessment of the current state of the GP service in Northland (New Zealand) and its accessibility. The second step applies a spatial optimisation framework to determine ideal spatial configurations that balance the trade-off between two conflicting objectives: 1) maximisation of accessibility and 2) the minimisation of infrastructure costs.

Maximisation of accessibility is pursued by minimising travel times between patients' residences and GP clinics, while the minimisation of infrastructure costs is intended as minimisation of the number of clinics and GPs.

3.1. Preliminary analysis: Northland's current GP distribution and accessibility evaluation

Currently, 31 clinics are operating in Northland for a total of 90 GPs practising, their distribution is illustrated in Figure 2 (adjacent clinics are grouped to ease visibility). Combining this data with the 2013 census population by meshblock dataset, it is possible to evaluate the current average patients/GPs ratio for different locations and also the average travel times required to access GP clinics from patients' households (together with the 75% and 95% quantile); the same data are summarised in Table 1. The patients/GPs ratio is calculated basing on two main assumptions: 1) patients are assigned to the closest clinic to their household and 2) travel times are evaluated starting from meshblock centroids.

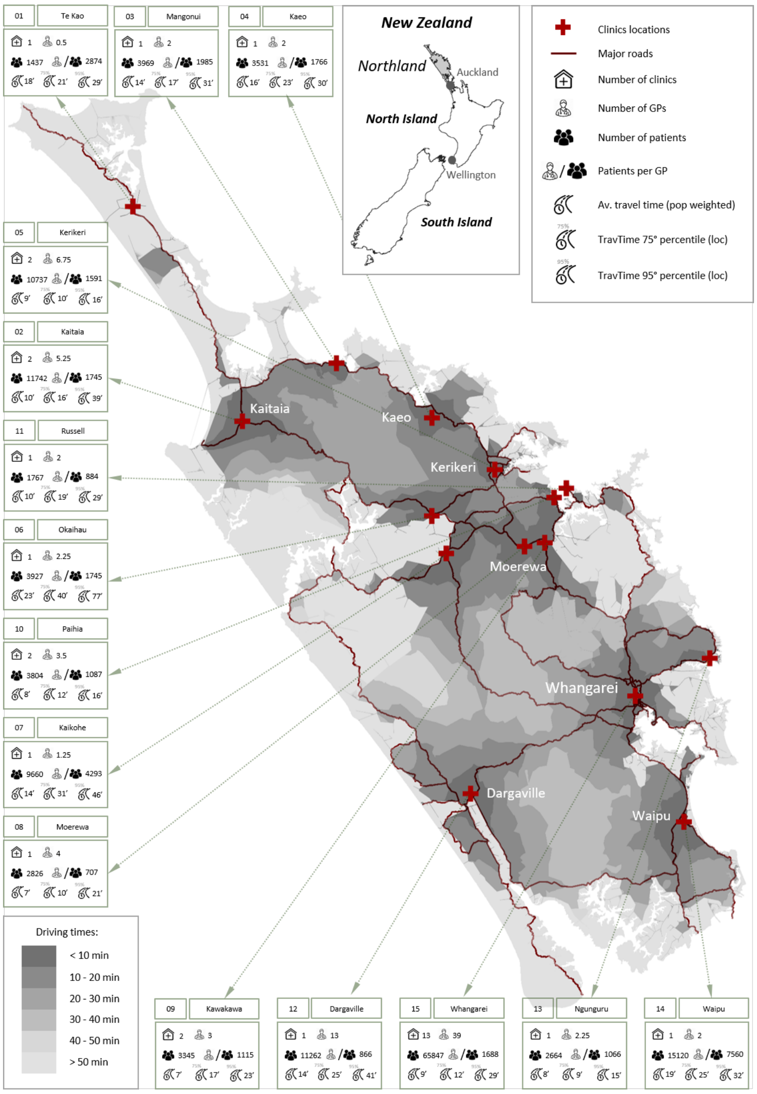

Figure 2. Map of Northland showing: number of clinics, number of GPs, population served, patients/doctors ratio, average, 75% and 95% quantile of travel times between households and clinics.

As it can be noted in both Table 1 and Figure 2, there is a significant variability in the patients/GPs ratio in different areas of the region, resulting in a considerable difference in the level of service provided. This is due to the rural nature of this region and the uneven distribution of services, facilities and infrastructure relative to how densely populated areas are.

Figure 4 shows the input data considered in this study; it shows Northland's road network and the Meshblock centroids from which distances and travel times are measured. In addition, an indication of all the available locations for the allocation of a GP clinic is provided; every town in the region is considered a potential available site for a clinic in the optimisation framework.

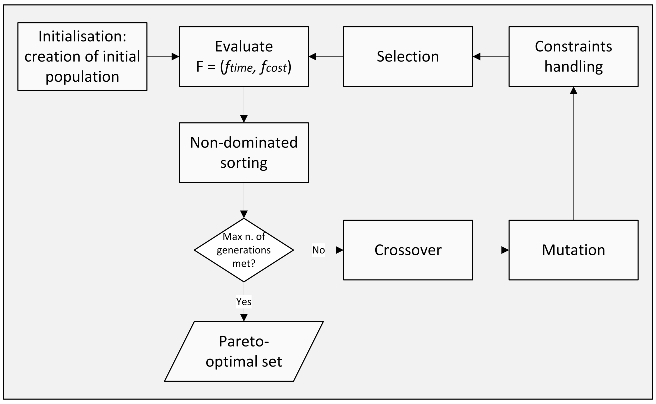

Figure 3. Genetic Algorithm flowchart.

Figure 4. Input data: Northland's road network, Meshblocks' centroids and available locations for GP clinics.

3.2. Spatial optimisation framework

After the assessment of the current situation of GP accessibility in Northland, the second phase of our analysis adopts a greenfield approach aimed at determining ideal distributions of clinics in the region that minimise travel times and minimise the infrastructure costs associated with GPs and clinics.

A heuristic multi-objective spatial optimisation framework that makes use of a genetic algorithm (GA) has been designed and applied to heuristically determine Pareto-optimal spatial plans that provide clinic configurations balancing the trade-off between two conflicting objectives (Figure 3). The first objective is the minimisation of travel times between patients' households and GP clinics; the second objective is the minimisation of the infrastructure costs (a function of the the number of clinics and GPs).

The framework consists in an iterative process where the number of generations G is an input parameter defined by the user; if the user does not have a prior knowledge of the case study, it is possible to calibrate the model by varying the number G to determine which is the most appropriate value. Low values of G allow quick runs, but they do not guarantee exhaustive searches of the solution space, while high values of G allow the inspection of a wider share of the solutions space with higher computational costs, and therefore higher run times. An effective calibration technique consists in analysing the evolution of the Pareto front. Typically, after a certain number of generations (50 in this application), the evolution of Pareto fronts slows down and after a number G of generations, there is no further significant improvement in the quality of solutions as shown in Figure 5.

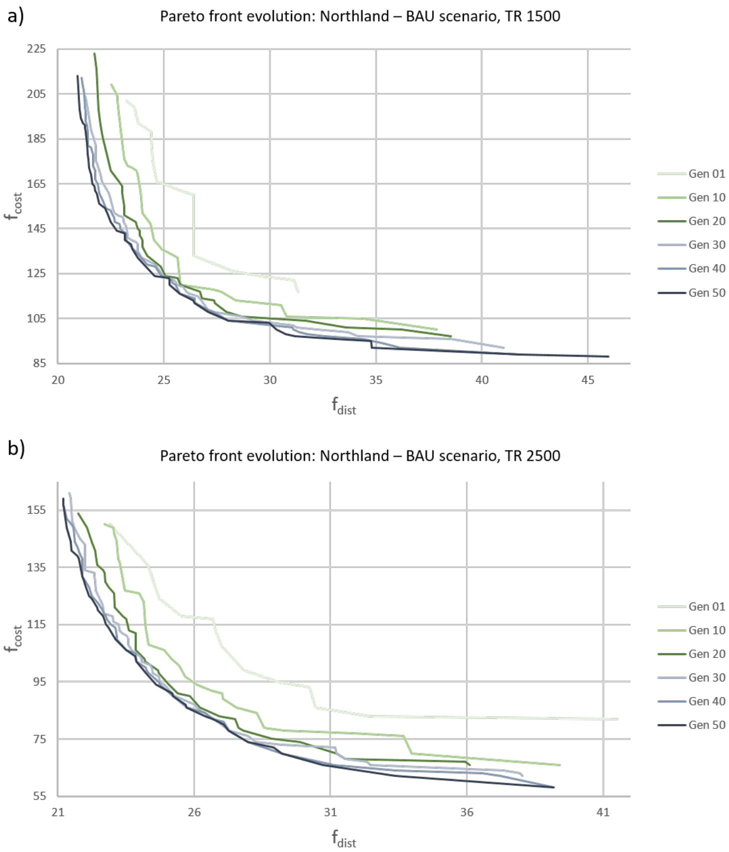

Figure 5. Pareto fronts evolution of the Northland case study: a) Target ratio patients/GPs 1500; b) Target ratio 2500.

Each generation has a parent set of spatial plans ($parents_g$) from which the child set ($offspring_g$) is generated by the application of the three evolutionary operators: *selection*, *crossover* and *mutation*. In the following iteration, after the application of the *selection* operator, a new set of parents is created: $parents_{\mathrm{g+1}}$. At this point, the process is repeated assigning $parents_g = parents_{\mathrm{g+1}}$.

Following this logic, the first generation of solutions (Generation 0) is created starting from the spatial plans generated in the initialisation phase combining a different number of locations and of clinics at these locations and different numbers of GPs assigned to each clinic. This set of initial potential solutions is evaluated against the two objectives, the poorly performing solutions are discarded while the best solutions are selected, combined and mutated to create a new generation of solutions (Generation 1). The final generation ($offspring_G$) presents Pareto-optimal solutions evolved from Generation 0.

Pareto-optimality is based on the concept of *domination*, which determines the best solutions in the balancing of the trade-off between the two conflicting objectives. If a solution A *dominates* a solution B, it means that A is not worse than B in all the objectives and it is strictly better in at least one. In the GA algorithm, if A dominates B, A will survive in the evolution process and B will be discarded. The objective values of solutions for the different conflicting objectives of this problem are determined according to the mathematical formulation of the two objectives. Given a current set of solutions, one of its members is called "Pareto-optimal" (within this set) if it is not dominated by another solution in the set. Our GA heuristically determines a set of Pareto-optimal solutions among the solutions considered in its iterations. The set of objective vectors of Pareto-optimal solutions is also known as the Pareto-front.

The adopted selection methodology is the Non-dominated Sorting Genetic Algorithm II (NSGA-II). NSGA-II is particularly appropriate for the solution of this kind of problems and it is extensively adopted (Cao et al. (2011); Caparros-Midwood, Barr, and Dawson (2017); Elkabalawy and Al-Sakkaf (2021); Jaeggi, Parks, Kipouros, and Clarkson (2008); W. Zhang and Fujimura (2010)). W. Zhang and Fujimura (2010). NSGA-II is more efficient for multi-objective optimisation (compared to other widely utilised algorithms, like the Strength Pareto Evolutionary Algorithm) in the estimation of the Pareto front. Differently from other traditional optimisation techniques (like simulated annealing and tabu search), Genetic Algorithms are able to explore broader areas of the solution space to avoid local optima (Yin, Xiao, Wen, and Fang (2017)) This is achieved combining the NSGA-II *selection* technique to the other evolutionary operators: the *crossover* operator is applied with a probability $p_{\mathrm{crossover}} = 0.6$, and it works by cutting two solutions S^1 and S^2 in two points cx_1 and cx_2 randomly chosen such that $0 < cx_1 < cx_2 < L$ (with L = length of S^1 and S^2). Their attributes are swapped in the central part of the list (i.e. between cx_1 and cx_2) and two new solutions $\mathrm{S}^{1'}$ and $\mathrm{S}^{2'}$ are created. Subsequently, with a probability $p_{\mathrm{mutation}} = 0.3$, the *mutation* operator is applied to those solutions on which the crossover has not been applied. In this process, a mutation of a randomly selected *i,j* location transforms the solution S into the new solution S'. The combination of these different operators has two advantages: the possibility to improve the performance of a solution in one or more objectives and the prevention of convergence on a small subset (local optimum) introducing new random locations and widening the search area in the solutions' space. The parameters were chosen as the result of a calibration and based on the work of Caparros-Midwood et al. (2017) in order to balance the trade-off between searching a large enough solution space and speeding up the search.

A number of necessary constraints are implemented in the planning framework, such as the target patients/GPs ratios that govern the allocation of GPs in different clinics, together with the maximum and minimum number of GPs allowed in each clinic and in the entire region. The constraints handling module is designed to evaluate each spatial plan and to discard potential solutions that do not meet the constraint. The target ratio defines both the upper and lower bounds of the allowed number of doctors for both the whole region and for each available location for clinics allocation. The constraint on the total number of practitioners in Northland is defined by setting a minimum and a maximum number of doctors that are obtained by multiplying the total population by the target ratio with a tolerance of ±25%. Again, the value was chosen as a result of a calibration, limiting the variability while preventing a too narrow solution space. Of course, the parameter is easily modifiable as the algorithm itself is customisable, if a user wanted to explore different tolerance parameters. An additional function in the constraints handling module defines the minimum and maximum allowed number of doctors for each available location for clinics allocation: this is achieved by multiplying the target ratio by the residing population within a 20 minutes radius from that location with a tolerance of ±50%.

In addition, other constraints are implicitly taken into account in the geographical data definition and in the evaluation of distances. They consist in the fact that households are aggregated at the meshblock level and travel times are evaluated from meshblocks' centroids. Moreover, travel times are calculated based on the road type. Different road types determine different allowed speeds and these speeds are the base for the calculation of travel times. This methodology has two main limitations. First, not all the roads of the same typology have the same allowed speed; this can be solved in future developments of this work with better input data (not available for the presented results). Also, this methodology evaluates a best case scenario as travel times are evaluated assuming free flow speed; traffic is not considered as a variable of the problem as well as the fact that some further speed limitations might be necessary due to roads layout (e.g. it may not be possible to always travel at the speed limit on very winding roads). This is a potential improvement to implement in future development of the present work.

3.2.1. Objective 1: Minimisation of travel cost

Travel times are evaluated as a result of a network analysis performed on the road network dataset available on the LINZ (Land Information New Zealand) Data Service of the New Zealand Government. In every potential solution (spatial plan), we calculate a weighted average travel time from each meshblock centroid to the closest clinic. The weighting factor consists of the number of served patients, so longer distances are penalised directly proportionally to how many people live far away from the GP clinics.

$$f_{time} = \frac{\sum_{m=1}^{M} TT_m \cdot Pop_m}{\sum_{m=1}^{M} Pop_m} \tag{1}$$

Where:

- f_{time}: travel time cost fitness of Objective 1;
- M: total number of meshblocks;
- TT_m: travel time from the centroid of meshblock m to the closest clinic;
- Pop_m: population living in meshblock m.

3.2.2. Objective 2: Minimisation of infrastructure costs

The second objective is the infrastructure cost function. We assume the cost is directly proportional to the number of clinics and practising GPs. Spatial plans with a high number of clinics and GPs are penalised with respect to this cost function. The mathematical formulation of f_{cost} consists of a weighted sum of the numbers of clinics and the number of GPs present in each spatial plan.

$$f_{cost} = \gamma_C \cdot n_C + \gamma_D \cdot n_D \tag{2}$$

Where:

- f_{cost}: infrastructure cost fitness of Objective 2;
- γ_C: weighting factor for clinics;
- n_C: total number of clinics of the evaluated solution;
- γ_D: weighting factor for GPs;
- n_D: total number of GPs of the evaluated solution.

The weighting factor for GPs γ_D has been assumed equal to 1 and the weighting factor for clinics γ_C has been assumed equal to 4. In general, this means that - with this formulation - it is not worth allocating a new clinic if the number of extra patients is smaller than 4 times the optimal patients/GPs ratio. For example, when the target ratio (TR) is 1 GP per 1500 patients, in a given spatial plan the algorithm will penalise spatial plans with clinics with a panel size of less than 6000 patients. Instead it will be considered more cost-effective to add an additional GP to an already allocated clinic nearby. Of course, the optimal location balances this against the travel time for those patients that is evaluated by the other objective function, f_{time}. The value of 4 was chosen as the number of full-time GPs in a medium-sized practice in New Zealand (Goodyear-Smith and Janes (2008); Leitch et al. (2018)). This is comparable to an average size practice in the United Kingdom (Kelly and Stoye (2014)) and a small-sized practice in the United States (Casalino, Devers, Lake, Reed, and Stoddard (2003)). Note that in this work, we define panel size as the number of patients that are allocated to a practice. The target ratio then expresses the expected number of patients per GP.

In addition to the choice of $\gamma_C = 4$, other values have been explored in this analysis as γ_C is a relevant model parameter and changing its value has consequences on the results. The choice of the value of γ_C can be considered as a policy scenario, as different values correspond to different aggregation/distribution of clinics on the territory. Depending on how costly / convenient opening a new clinic can be, different values of γ_C might be more appropriate. Since the presented framework is meant as an exploratory tool, we present a sensitivity analysis of this parameter (see following sections) to assess its impact on final Pareto-optimal spatial plans.

3.3. Mathematical model

In order to further clarify the underlying location problem, a corresponding mathematical model is formulated. It uses the following notation:

- I: set of clinic locations
- J: set of demand locations, i.e. meshblocks, with $|J| = M$;
- t_ij: travel time from clinic i to the centroid of meshblock j;
- p_j: population living in meshblock j.

- R: ratio patients per GP;
- C: weighting factor for clinics;
- D: weighting factor for GPs;
- u_i: upper bound on the number of GPs for clinic $i \in I$;
- l_i: lower bound on the number of GPs for clinic $i \in I$;
- U: upper bound on the number of GPs in total;
- L: lower bound on the number of GPs in total;
- S: upper bound on exceeding the patients-per-GP-ratio;
- $y_i \in \{0;1\} \forall i \in I$: 1, if a clinic in location i is opened, 0 otherwise;
- $x_{ij} \in \{0;1\} \forall i \in I, j \in J$: 1, if a clinic in location i covers patients in meshblock j, 0 otherwise;
- $z_i \in \mathbb{N}_0 \forall i \in I$: number of GPs working in clinic i.

The model can then be formulated as follows:

$$f_{time}: \text{ minimise } \frac{\sum_{i,j} t_{ij} \cdot p_j \cdot x_{ij}}{\sum_j p_j} \tag{3}$$

$$f_{cost}: \text{ minimise } \sum_i (C \cdot y_i + D \cdot z_i + E \cdot (s_i + m_i)) \tag{4}$$

subject to

$$x_{ij} \leq y_i \qquad \forall i \in I, j \in J \tag{5}$$

$$\sum_{i \in I} x_{ij} = 1 \qquad \forall j \in J \tag{6}$$

$$\sum_{i \in i} z_i \geq L \tag{7}$$

$$\sum_{i \in i} z_i \leq U \tag{8}$$

$$z_i \geq l_i \cdot y_i \qquad \forall i \in I \tag{9}$$

$$z_i \leq u_i \cdot y_i \qquad \forall i \in I \tag{10}$$

$$\sum_{j \in J} x_{ij} \cdot p_j = R \cdot z_i - s_i + m_i \qquad \forall i \in I \tag{11}$$

$$s_i = 0 \vee m_i = 0 \qquad \forall i \in I \tag{12}$$

$$s_i <= S \cdot z_i \qquad \forall i \in I \tag{13}$$

$$y_i, x_{ij} \in \{0;1\} \qquad \forall i \in I, j \in J \tag{14}$$

$$z_i \in \mathbb{N}_0 \qquad \forall i \in I \tag{15}$$

$$s_i, m_i \in \mathbb{R} \qquad \forall i \in I \tag{16}$$

The objective functions (3) and (4) are similar to those described above. While (3) is equivalent to (1), (4) needs to be slightly adapted to account for the potential deviations from the target ratios that are to be minimised as well. Constraints (5) ensure that a meshblock can only be served by an opened clinic. Each meshblock must be served and assigned to a clinic location (6). Constraints (7) to (10) address the upper and lower bounds on the numbers of GPs per clinic and in total. In addition, constraints (11) determine the surplus or malus of assigned patients to a clinic in relation to the desired patients per doctor ratio. Constraints (12) ensure that either the target ratio is exceeded or undercut, but prevents both decision variables to be strictly greater than zero. The equation can also be linearised, the form was chosen for brevity. Note that it could also be solved by the CPLEX solver in this form, for example, as CPLEX is able to linearise it easily. The ratio cannot be exceeded by S patients per doctor in an opened clinic (13). Constraints (14) to (16) are the domain constraints.

The model was implemented in the IBM ILOG Optimization Studio and solved with CPLEX on an AMD Ryzen 7 Microsoft Surface (R) Edition laptop with 16 GB RAM. A weighted sum approach was used. In future work, a dichotomic approach could be applied to determine a set of Pareto efficient solutions for the problem Raith and Ehrgott (2009). Due to the very long run times and our focus on the genetic algorithm framework and the resulting decision support tool, it was omitted in this work. Solving already one instance with one set of weights and one target ratio (1500) exceeded a run time of 120 minutes, when the run was stopped. After a run time of 60 minutes, the gap was still at 13.46 %.

3.4. Genetic algorithm results

The methodology and the genetic algorithm framework described in Section 3.2 has been applied to the Northland instance with two different TRs between patients and GPs to explore and compare different scenarios. The chosen TRs are 1500 and 2500 as reasonable average patient numbers for GP practices in rural areas in New Zealand to ensure sufficient income and avoid overcrowding (Goodyear-Smith and Janes (2008)). Note that efficient panel sizes and therefore target ratios might vary for different regions or even within a region and for some regions lower panel sizes might be beneficial (Altschuler, Margolius, Bodenheimer, and Grumbach (2012); Dahrouge et al. (2016)). Therefore, a thorough analysis only on the panel size and target ratio might be necessary, which will be further discussed in the outlook.

Figure 6 shows all the inspected solutions and the Pareto-fronts in the solution space defined by the two optimisation functions: f_{time} (minimisation of travel times) and f_{cost} (minimisation of infrastructure costs).

The horizontal axis represents the weighted average of travel times (as captured by f_{time}, where weights represent the served population) and it is expressed in minutes; while the vertical axis represent the weighted sum of the number of clinics and GPs of the different solutions. The blue points represent all the potential solutions inspected by the genetic algorithm and the ones lying on the Pareto-front (indicated by the points that lie on the dashed red line) are the ones that optimally balance the trade-off between the two conflicting objectives.

Every blue point in Figure 6 represents a spatial plan of clinics in different locations, each with a different number of GPs. In Figures 7, 8, 9 and 10, a few examples are provided to illustrate different solutions lying in different areas of the Pareto-front for both the considered TRs. For each Pareto-front, two solutions are selected (Figure 6): one from the top-left area of the Pareto-front (solutions A and C: low travel times, high number of clinics and GPs) and one from the bottom-right area (solutions B and D: lower infrastructure costs, but higher average travel times).

Solutions A and C correspond to the same f_{time} fitness, as well as solutions B and D. These two pairs of solutions are chosen as they represent two thresholds beyond which there is still an improvement in one of the two objectives, without a significant improvement in the other one. In fact, beyond A and C there is an increase of infrastructure costs without a significant decrease in travel times. Similarly, beyond B and D there is an increase in travel times, without a significant decrease of infrastructure costs.

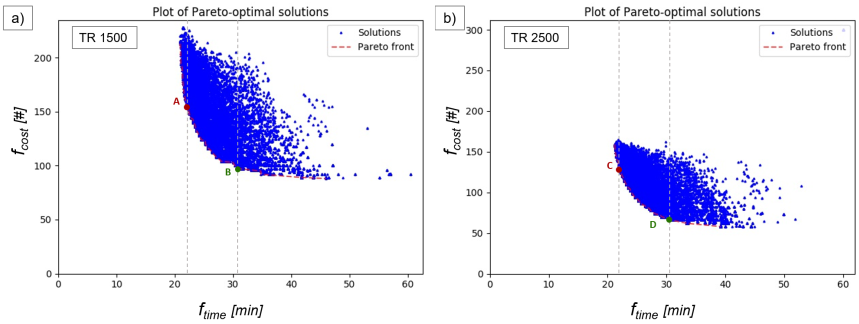

Figure 6. Solutions and Pareto-fronts for the two evaluated scenarios: a) TR 1500 and b) TR 2500, together with the extracted solutions (A, B, C and D) represented in Figures 7, 8, 9 and 10.

The maps in Figures 7, 8, 9 and 10 provide several insights into the Pareto-optimal spatial plans, among which there is the indication of the number of allocated patients to each clinic of the solution and the number of patients residing in a 20 minute radius (catchment area) from the clinic itself. Each clinic location has an allowed range of potential number of GPs (in square brackets) that is proportional to the population residing within the catchment area. This feature provides several benefits to the optimisation procedure: first, it speeds up running times as it constrains the number of possible solutions. Then, it also mitigates another assumption of this methodology: the fact that patients are allocated to the closest clinic. In remote areas (like the North of Northland), this is a sensible assumption, but in the southern part of this region, where travel times are lower and more similar to one another, considering a 20 minute catchment area allows potential competition of close by clinics to be taken into account. In fact, patients can freely choose a clinic and if several are within the same distance, these could be seen as competing with each other. A recent study has determined a driving time of 15 minutes to the closes GP to be widely accepted in rural Germany (Schröder, Flägel, Goetz, and Steinhäuser (2018)). As mentioned before, the population density in rural Germany strongly exceeds the population density in rural New Zealand. Therefore, we assume a 20 minute radius to be acceptable in Northland, New Zealand.

In the lower left corner of each figure, a graph shows the patients per GP ratio for each opened location together with lines for the target ratio and the regional average. Due to the high variation in population numbers and the long distances, it is difficult to even the patient ratios for all locations. While in Figures 7 and 9 the regional average is below the target ratio, it exceeds it in Figures 8 and 10. These graphs could therefore provide valuable insights for decision makers when comparing different possible GP location.

As anticipated in the previous section, in addition to this analysis, several values of

γ_C, the weighting factor for clinics in f_{cost}, are inspected (namely $\gamma_C = 2$, 4 and 6) to assess how results vary when we change this parameter. To complete the sensitivity analysis, also a wider set of ratios has been explored (1000, 1500, 2000, 2500 and 3000). The result of this is a set of 15 different scenarios; each one of them has its own Pareto-front, which, on average, consist of 35-50 Pareto-optimal spatial plans, for a total of 595 solutions.

Among these 595 solutions, we analysed those with the same accessibility f_{time} as the solutions presented in Figures 7, 8, 9 and 10; i.e. we present a comparison of the infrastructure costs (f_{cost}) for fixed values of f_{time} ($f_{time} = 22$ min and $f_{time} = 31$ min); these values are derived from the thresholds shown in Figure 6 (i.e. the f_{cost} and f_{time} values of solutions A, B, C and D). This comparison is summarised in Table A1 (Appendix).

4. Discussion

The results presented in the previous section are illustrative as many input parameters can be changed to evaluate different scenarios. The optimisation framework is flexible in this regards, as versatility has been one of the main design criteria in its development. Among others, the main parameters that the final user can vary to evaluate different scenarios are:

- The TR between patients and GPs. (The spatial plans in the results section explore the target patients/GPs ratio of 1500 and 2500. Table A1 in the appendix also shows the results for TR = 1000, 2000 and 3000.)
- The value of γ_C (weight coefficient for clinics in the cost function). The spatial plans in the result section have a fixed value of value of $\gamma_C = 4$; Table A1 in the appendix also shows the results for $\gamma_C = 2$ and 6.
- The road network: we propose an analysis with a business as usual road network, but several scenarios can be tested – for instance, an assessment of primary care accessibility during road closures due to disruptions caused by natural events (such as floods or landslides).
- The cost associated with GPs and clinics. Here we assume the infrastructure costs are directly proportional to the number of GPs and clinics, but more sophisticated cost functions can be implemented if better data are available – e.g. land values or rent costs.
- The average travel time to designate clinics catchment areas. In the proposed results, catchment areas of 20 minutes' driving time have been considered in relation to previous research published for rural Germany (Schröder et al. (2018)).

Given these assumptions, Pareto-optimal spatial plans like the ones shown in the previous section can help local authorities to assess the healthcare service level and the impacts of potential interventions. Exploring different TRs between GPs and patients allows the cost-effectiveness of opening new clinics in different areas to be explored. For example, ratios that are too high would imply many patients per GP (consequently, a profitable business for the practitioner). This, in turn, may result in long waiting lists and ultimately a lower accessibility to the healthcare service. On the other hand, low ratios would imply an abundance of GPs, with the potential risks of competition being too high, resulting in the lowering of fares and struggles in maintaining a profitable business (that may even lead to potential closures).

The presence of these conflicting objectives on different levels (physical and eco-

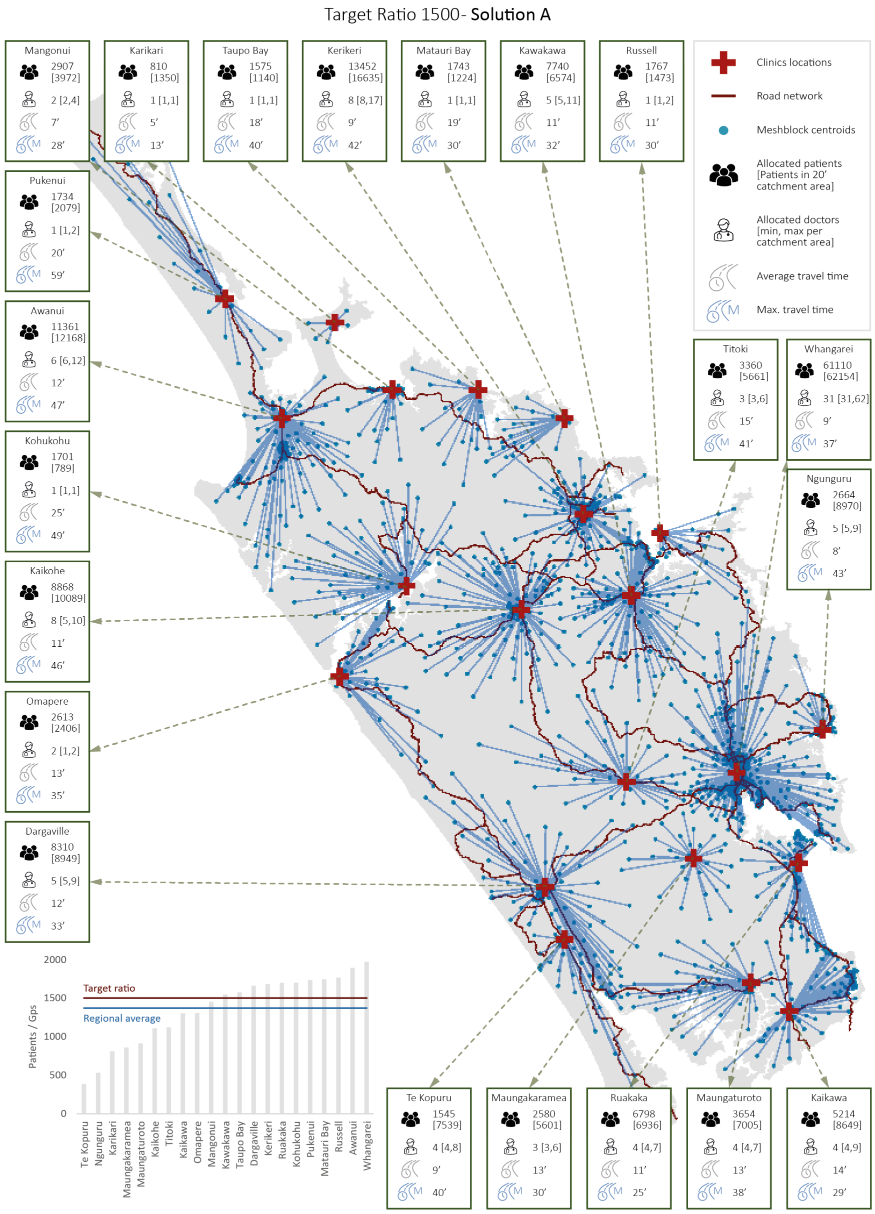

Figure 7. Target patients/GPs ratio: 1500, Solution A. The blue straight lines indicate the allocation of the different meshblocks to the closest clinic.

nomic) makes these considerations challenging without any auxiliary tool able to provide quantitative evidence supporting the different explored scenarios. The proposed framework can help make to support better and more informed decision making on resource allocation and policy making by local authorities.

Figure 11 provides an example of how the proposed results are useful in comparing different spatial plans and different scenarios. The spatial variability between different solutions may seem low when comparing similar solutions; this is why the maps in Figures 7 - 10 provide also quantitative values attributed to the different locations (i.e. statistics regarding the served population, the allocation of GPs, travel times etc.). Since solutions A and C represent two solutions with the same travel cost fitness f_{time}, the two plans are spatially very similar (this happens also to solutions B and

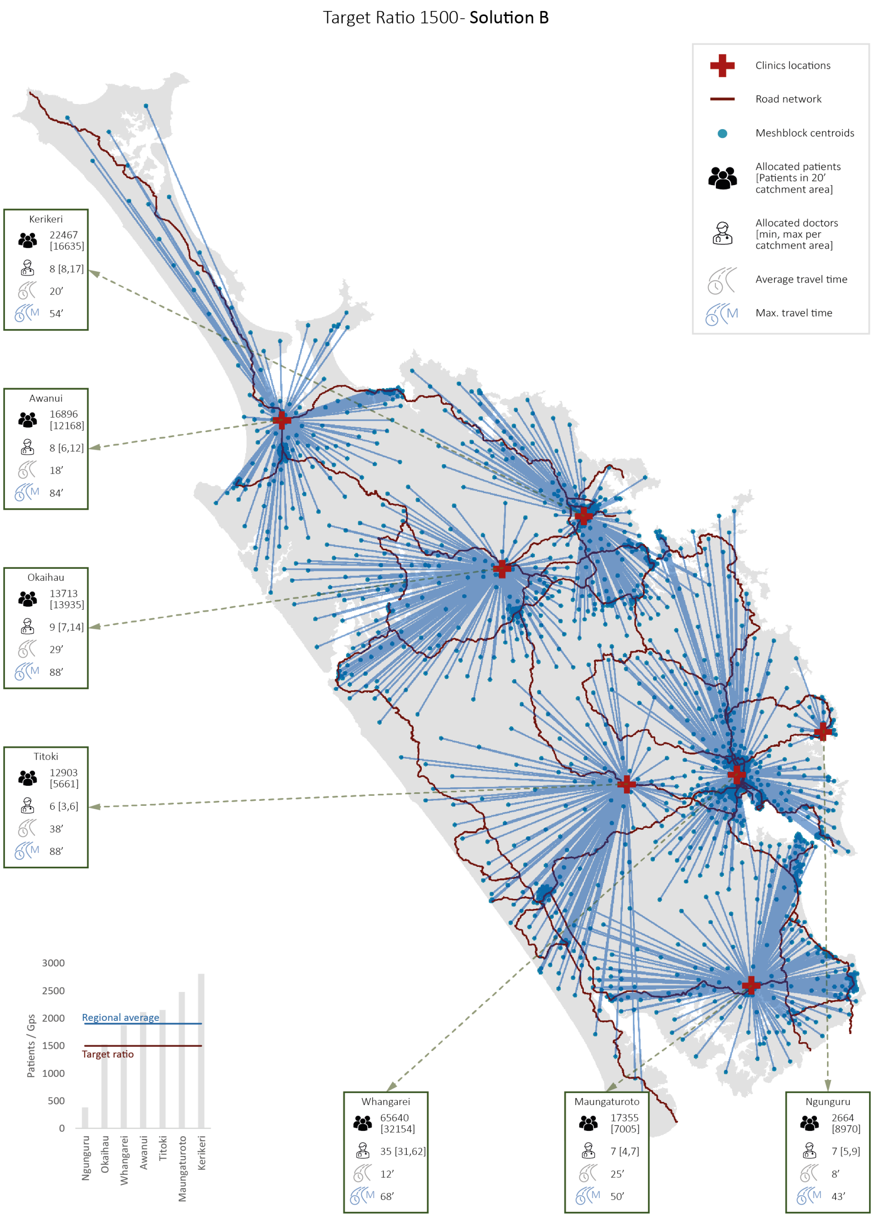

Figure 8. Target patients/GPs ratio: 1500, Solution B. The blue straight lines indicate the allocation of the different meshblocks to the closest clinic.

D). However,they differ from the infrastructure cost perspective: Solution A involves 100 GPs (+11% with respect to the current level of employment), while Solution C involves 61 GPs (-32%). Solution C provides the same accessibility at a lower infrastructure cost, nevertheless, Solution C is not more cost effective tout court as this improvement comes at a cost: a higher TR. When comparing different solutions from different scenarios, maps like those shown in Figure 11 provide an immediate sense of how accessibility varies when changing different parameters (like the TR and the position in the Pareto-front).

A useful comparison is also possible with the current distribution of clinic locations and GPs in Northland (Figure 2). The current service is considered as a baseline to assess the improvements provided by the optimisation framework's solutions. Table 2

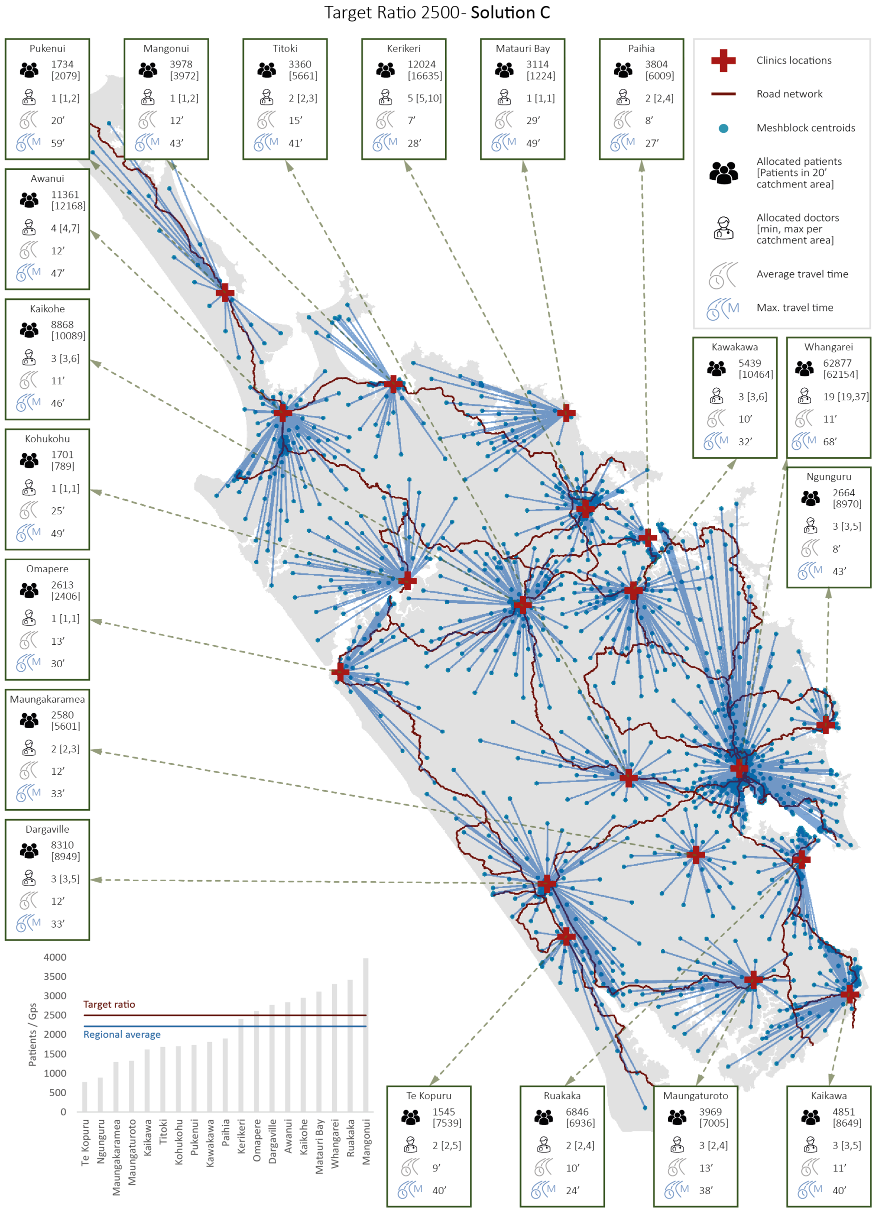

Figure 9. Target patients/GPs ratio: 2500, Solution C. The blue straight lines indicate the allocation of the different meshblocks to the closest clinic.

provides an overview of this comparison in terms of relative percentage increments (or reductions) of served population with different travel times. Solutions A and C present improvements in accessibility (i.e. more people served in less than 20 minutes). This is achieved with an increase of 11% in the number of GPs when considering a target patients/GPs ratio of 1500, and with a reduction of almost a third of the number of GPs when considering a ratio of 2500. Solutions B and D, instead, are those spatial plans from the bottom-right area of the Pareto-fronts. This implies lower accessibility and lower costs. These solutions are useful to understand the consequences of potential reductions in the number of available GPs. According to the results analysis summarised in Table 2, they necessarily imply reductions in the served population in less than 20 minutes, but provide useful information regarding optimal allocations of

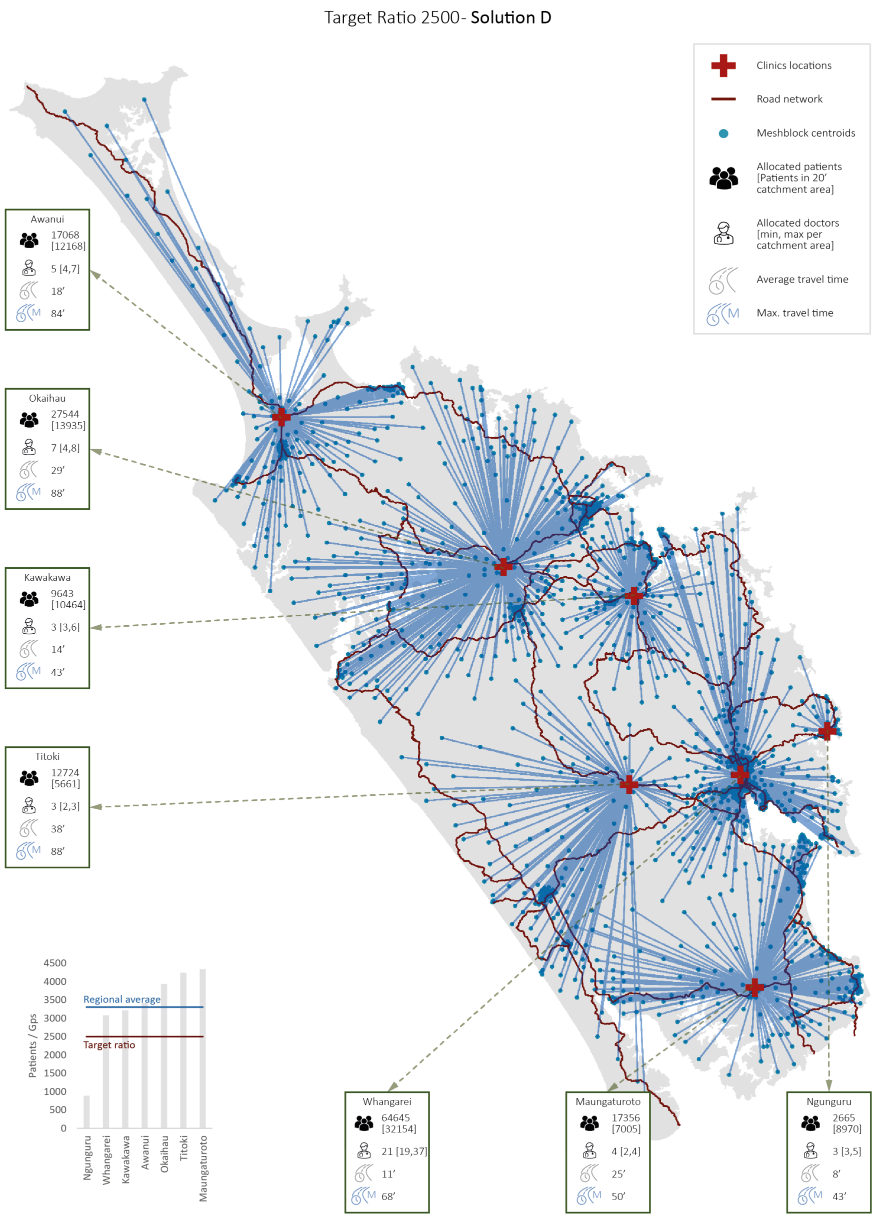

Figure 10. Target patients/GPs ratio: 2500, Solution D. The blue straight lines indicate the allocation of the different meshblocks to the closest clinic.

a reduced number of resources.

Parameters $\gamma_C = 4$ and TR $= 1500$ and 2500 have been chosen as the most sensible choices in the presentation of the results of the spatial optimisation framework. The choice of the TR has been explained in the previous sections, nevertheless an additional set of TRs (1000, 2000 and 3000) has been explored to understand the sensitivity of the system with respect to this parameter. Also the choice of $\gamma_C = 4$ has been justified in the definition of the infrastructure cost function, however, we wanted to test the sensitivity of the system also to this parameter. The results of this analysis is summarised in Table A1 (in the Appendix).

Since in both tables, Table 2 and Table A1, we are comparing different spatial plans with the same f_{time} fitness, it is not surprising that across the different values of γ_C

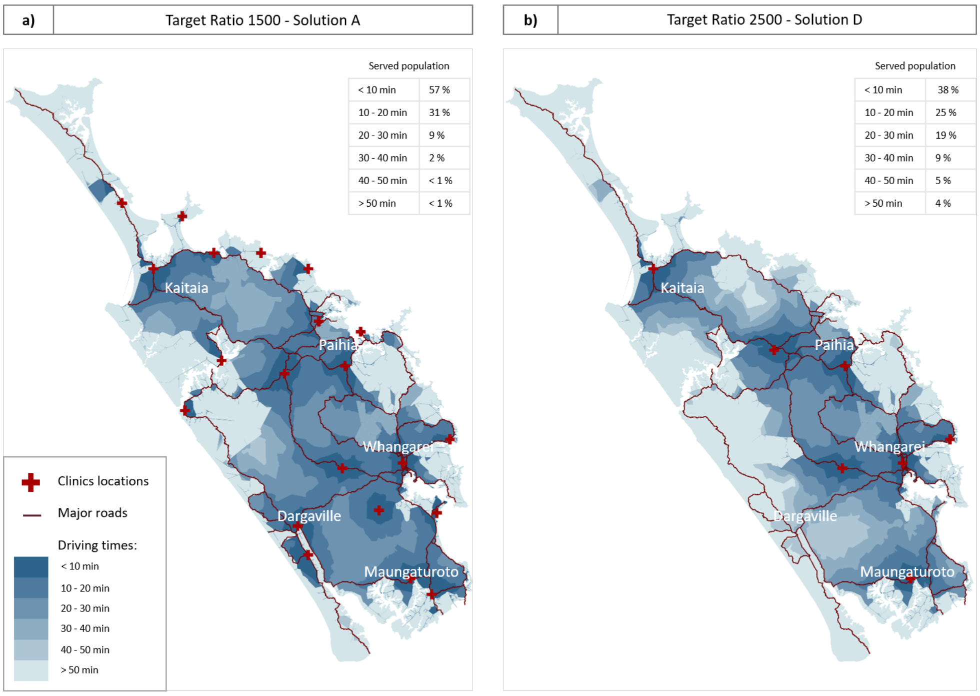

Figure 11. Accessibility comparison between different Pareto-optimal spatial plans in different scenarios. a) Solution A (target patients/GPs ratio: 1500). b) Solution D (target patients/GPs ratio: 2500).

and TR, only small differences can be observed from the accessibility perspective: +2 to +4% of the population served under 10 minutes with respect to the baseline (i.e. current situation in Northland) for $f_{time} = 22$ min. However, it is interesting to observe how the infrastructure cost varies (i.e. number of FTE GPs and clinics) to achieve such an improvement in terms of accessibility (i.e. % of population served under 10, 20, 30 etc. minutes).

Figure A1 (in the Appendix) shows that in all the spatial plans it is necessary to increase the current number of practising GPs when considering a TR of 1000. Allowing more patients per GP (i.e. increasing the TR) it is possible to achieve the aforementioned improvement in terms of service accessibility with a lower number of GPs compared to the current number of practising GPs in Northland (i.e. 90). Also, different values of γ_C correspond to different numbers of employed GPs; this happens because of the mathematical formulation of the infrastructure cost function: lower values of γ_C correspond to lower costs in opening new clinics compared to adding GPs to already open clinics. This parameter is not directly proportional to the infrastructure cost fitness as it is intended as a parameter to model the competition between close-by available locations for clinics. However, we can observe that for all the values of γ_C, the total number of GPs in the region decreases when TR increases: varying γ_C's value does not affect the behaviour of the solutions when varying TR. Consequently, we consider TR as the governing parameter when comparing travel cost and infrastructure cost fitnesses and γ_C as a secondary modelling parameter to capture local competition.

Understanding the sensitivity of the system to the governing parameters and the main assumptions is useful to calibrate the model, represent the current situation and discuss potential improvements to the in-place healthcare infrastructure, but also for modelling future scenarios. For example, if the population pattern stays the same, the

Table 2. Comparison between Northland's current GP practices accessibility and the optimisation framework's results.

Travel time		<10 min	10-20 min	20-30 min	30-40 min	40-50 min	>50 min	
Current situation	Pop served	83277	39423	18750	7554	1557	1077	**N of GPs**
	% Pop	**55**	**26**	**12**	**5**	**1**	**1**	90
TR 1500 Sol A	Pop served	86385	46905	14094	3513	657	84	**N of GPs**
	% Pop	57	31	9	2	0	0	100
	±%	**+2**	**+5**	**-3**	**-3**	**-1**	**-1**	+11%
TR 1500 Sol B	Pop served	61791	36840	24105	16512	6768	5622	**N of GPs**
	% Pop	41	24	16	11	4	4	80
	±%	**-14**	**-2**	**+4**	**+6**	**+3**	**+3**	-11%
TR 2500 Sol C	Pop served	88221	42948	13494	3558	1818	1599	**N of GPs**
	% Pop	58	28	9	2	1	1	61
	±%	**+3**	**+2**	**-3**	**-3**	**+0.2**	**+0.3**	-32%
TR 2500 Sol D	Pop served	57424	38001	28314	14232	7026	6648	**N of GPs**
	% Pop	38	25	19	9	5	4	46
	±%	**-17**	**-1**	**+6**	**+4**	**+4**	**+4**	-49%

Tot Pop
151638

optimisation results remain valid, but future developments in Northland's demography might invalidate any conclusion drawn from the current situation; if some smaller towns grow faster than larger ones, if certain areas develop (or shrink) at considerably different rates, the spatial optimisation will have to be updated with respect to the input parameters.

Finally, as the framework is flexible, different rules and probabilities for allocating patients to practice locations can be explored to model how patients might choose a GP. Besides, in addition to travel time, costs for a GP visit can be a relevant factor in Northland when choosing a GP. As stated above, GP practices in New Zealand can define their costs individually, resulting in significant differences in fees between practices. In order to reasonably integrate such costs into the framework, more information on their influence on GP choice is necessary. If patient enrolment information is accessible, it can be used to study how patients choose a GP practice as part of future research. Another potential idea is to use deprivation indices Exeter, Zhao, Crengle, Lee, and Browne (2017) and average income values within meshblocks together with cost levels of GPs. Future research should investigate the catchment areas of a practice taking potentially all relevant factors into account, including distance and costs.

5. Conclusion

We have presented a genetic algorithm framework to locate primary care practices and analyse the accessibility and expected demand at current and new GP locations that can be used by managers or healthcare ministries. The framework was applied to the locations of primary care practices in Northland, New Zealand. Besides an analysis of the current GP practice locations, we have computed and discussed several alternative location structures. While these do not aim to suggest an immediate restructuring of the GP locations in Northland, they exemplify how the spatial optimisation framework can be used by decision-makers to balance multiple healthcare objectives and help decision makers identify least-bad tradeoffs.

The outputs of the framework presented in the previous sections have the potential to guide future investments in the healthcare sector in Northland. Table A1 (in the Appendix) provides indications on how to increase accessibility by investing in the infrastructure. For instance, for a fixed target patients/GPs ratio of 1500 and γ_C=4, if Northland wants to increase the population served in under 20 minutes by +7%, it should increase the number of GPs (in terms of FTE) by +11%. Different combinations of TR and γ_C can be explored according to the end-user's needs/interests.

In future research, the framework is to be extended by additional parameters and features as suggested in the previous section (e.g. including traffic, different cost levels for GP appointments etc.). In addition, more scenarios are to be investigated, such as considering potential disruptions of the road network, including different kinds of healthcare facilities (i.e. not only GP clinics) and applying the optimisation framework to different areas. A dedicated study should especially target efficient panel sizes, i.e. the number of patients per GP, which might vary within different Northland regions. In 2003, Scott et al. found that low-income groups and Māori were significantly less likely to visit a GP at least once in the year (Scott, Marwick, and Crampton (2003)). While this might serve as input, it could and should also be used to address the issue of equity in access for all inhabitants and investigate reasons that currently prevent equal access and care levels as well as countermeasures to overcome it. A prerequisite for this research is the availability of current data on GP visits per inhabitant and year throughout the region. Another idea could be to actually optimise the patient-GP-ratio within the framework, for example assuming a fixed budget and taking the expected income / workload as well as the expected availability into account in a bi- or multi-objective optimisation approach. Finally, it would also be interesting to explore a population change scenario, not only from a growth/shrinkage perspective, but also in terms of demographic change (e.g. considering different needs for different portions of population - like elderly people who might need more frequent care).

Another important topic for future research is the extension of the framework features and the visualisation of outputs to further support decision makers. For example, it might be interesting for them to access information for each location, e.g. by displaying a bar chart showing the outputs for the different panel or practice sizes. In order to estimate the importance of a GP practice at a certain location, we aim to develop an extension that finds those locations that are less sensitive to changes in the input parameters and that can compute the likelihood for a meshblock to have a GP practice for several different Pareto-optimal solutions. Often, decision makers only have a limited budget available to invest and are interested in especially attractive locations that can be accessed by many patients, for example. At the same time, not all input parameters might be fixed or different scenarios might exist for future developments. In their study, Reuter-Oppermann et al. (2019) have shown, for example, that a location that was optimal for the one scenario the decision maker was currently considering, was not included in any other future scenario. It will therefore be valuable to add a feature to the framework that produces outputs for single locations and supports an explicit analysis. As an alternative, future research could explore the development of spatial plans that slowly increase the number of clinic locations (or GPs) to identify the best sequence of developments to get to an optimal allocation.

Besides the aforementioned topics for future research, also the methodology used within the framework could be addressed. Probably, heuristics will still need to be part of the framework due to the unavailability of / high costs for mathematical solvers in practice as well as the very long run times that we experienced for the Northland case in this work. Nevertheless, we aim to apply the presented mathematical models to analyse the performance of the heuristic and to further improve it. In addition, exact solution approaches could be investigated. We further want to analyse until which problem size we might be able to apply the mathematical model.

The outputs of this framework should not be considered as prescriptive, rather the framework is an exploratory tool for use by decision makers. For example, District Health Boards and Primary Health Organisations can identify a promising subset of solutions that meet key criteria for which other qualitative factors can be considered in more detail. Application of this approach during periodic strategic reviews by health boards can thereby ensure more effective use of resources and healthcare coverage over time.

Appendix A.

Table A1 presents the summary of the accessibility analysis performed across all the inspected scenarios. Figure A1 presents a comparison of infrastructure cost compared to the current Northland situation for different values of γ_C and different TRs (for a fixed value of $f_{time} = 22$). It is worth noting that since this graph is heuristically computed, for TR=1500 and TR=2500, the case with γ_C=6 has fewer GPs than γ_C=4 and vice versa for other values of TR.

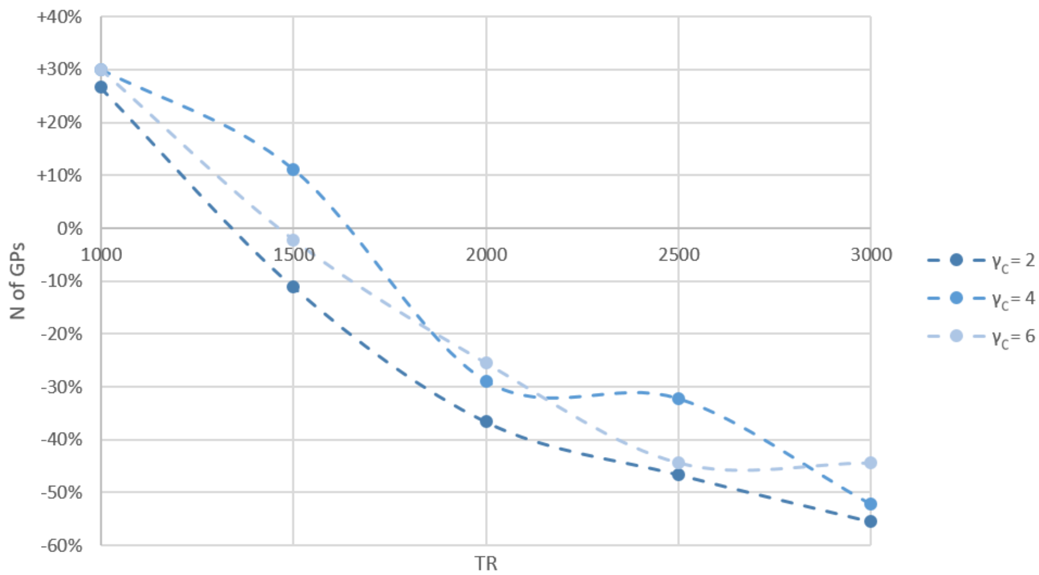

Figure A1. Comparison of infrastructure cost with respect to current Northland situation for different values of γ_C and different TR for a fixed value of $f_{time} = 22$.

References

Abernathy, W. J., & Hershey, J. C. (1972). A spatial-allocation model for regional health-services planning. *Operations Research*, *20*(3), 629–642.

Ahmadi-Javid, A., & Ramshe, N. (2020). A stochastic location model for designing primary healthcare networks integrated with workforce cross-training [Journal Article]. *Operations Research for Health Care*, *24*, 100226. Retrieved from `http://www.sciencedirect.com/science/article/pii/S2211692318300894`

Table A1. Accessibility analysis of all the inspected solutions for a fixed value of f_{time} = 22. Comparison between Northland's current GP practices accessibility and the optimisation framework's results.

			Travel time		<10 min	10-20 min	20-30 min	30-40 min	40-50 min	>50 min	
Current situation				**Pop served**	83277	39423	18750	7554	1557	1077	**N of GPs**
				% Pop	55	26	12	5	1	1	90
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86811	45114	13680	3543	891	1599	**N of GPs**
2	1000	49201	22	**% Pop**	57	30	9	2	1	1	114
				± %	+2	+4	-3	-3	-0.4	+0.3	+27%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86871	44943	15312	3744	684	84	**N of GPs**
2	1500	43040	22	**% Pop**	57	30	10	2	0	0	80
				± %	+2	+4	-2	-3	-1	-1	-11%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86232	45600	13044	3414	1749	1599	**N of GPs**
2	2000	39062	22	**% Pop**	57	30	9	2	1	1	57
				± %	+2	+4	-4	-3	+0.1	+0.3	-37%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	88248	45276	13899	3450	681	84	**N of GPs**
2	2500	49149	22	**% Pop**	58	30	9	2	0	0	48
				± %	+3	+4	-3	-3	-1	-1	-47%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	89175	43374	13866	3528	1611	84	**N of GPs**
2	3000	41995	22	**% Pop**	59	29	9	2	1	0	40
				± %	+4	+3	-3	-3	+0.04	-1	-56%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86973	45651	14079	4083	768	84	**N of GPs**
4	1000	47192	22	**% Pop**	57	30	9	3	1	0	117
				± %	+2	+4	-3	-2	-1	-1	+30%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86385	46905	14094	3513	657	84	**N of GPs**
4	1500	48254	22	**% Pop**	57	31	9	2	0	0	100
				± %	+2	+5	-3	-3	-1	-1	+11%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86988	45507	14433	3942	684	84	**N of GPs**
4	2000	42010	22	**% Pop**	57	30	10	3	0	0	64
				± %	+2	+4	-3	-2	-1	-1	-29%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	88221	42948	13494	3558	1818	1599	**N of GPs**
4	2500	41222	22	**% Pop**	58	28	9	2	1	1	61
				± %	+3	+2	-3	-3	+0.2	+0.3	-32%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	89268	43398	13524	3669	1695	84	**N of GPs**
4	3000	43489	22	**% Pop**	59	29	9	2	1	0	43
				± %	+4	+3	-3	-3	+0.1	-1	-52%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	87720	44886	13842	3486	1620	84	**N of GPs**
6	1000	48080	22	**% Pop**	58	30	9	2	1	0	117
				± %	+3	+4	-3	-3	+0.04	-1	+30%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86988	45507	14433	3942	684	84	**N of GPs**
6	1500	42003	22	**% Pop**	57	30	10	3	0	0	88
				± %	+2	+4	-3	-2	-1	-1	-2%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	87159	45594	14229	3873	624	159	**N of GPs**
6	2000	41211	22	**% Pop**	57	30	9	3	0	0	67
				± %	+3	+4	-3	-2	-1	-1	-26%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	86811	45114	13680	3543	891	1599	**N of GPs**
6	2500	46014	22	**% Pop**	57	30	9	2	1	1	50
				± %	+2	+4	-3	-3	-0.4	+0.3	-44%
γ_C	**TR**	**Sol ID**	f_{time}	**Pop served**	87894	44571	13464	3945	1680	84	**N of GPs**
6	3000	35209	22	**% Pop**	58	29	9	3	1	0	50
				± %	+3	+3	-3	-2	+0.1	-1	-44%

Altschuler, J., Margolius, D., Bodenheimer, T., & Grumbach, K. (2012). Estimating a reasonable patient panel size for primary care physicians with team-based task delegation. *The Annals of Family Medicine*, *10*(5), 396–400.

Baum, M. A., Bergwall, D. F., & Reeves, P. N. (1975). Planning health care delivery systems [Journal Article]. *Am J Public Health*, *65*(3), 272-9.

Brabyn, L., & Skelly, C. (n.d.). Geographical access to services, health (gash): modelling population access to new zealand public hospitals. In *proceedings: Thirteenth annual colloquium of the spatial information research centre.*

Brabyn, L., & Skelly, C. (2002). Modeling population access to new zealand public hospitals. *International journal of health geographics*, *1*(1), 1–9.

Bullen, N., Moon, G., & Jones, K. (1996). Defining localities for health planning: A GIS approach [Journal Article]. *Social Science and Medicine*, *42*(6), 801 - 816. Retrieved from http://www.sciencedirect.com/science/article/pii/0277953695001808

Butler, D. C., Petterson, S., Bazemore, A., & Douglas, K. A. (2010). Use of measures of socioeconomic deprivation in planning primary health care workforce and defining health care need in australia [Journal Article]. *Aust J Rural Health*, *18*(5), 199-204.

Cao, K., Batty, M., Huang, B., Liu, Y., Yu, L., & Chen, J. (2011). Spatial multi-objective land use optimization: extensions to the non-dominated sorting genetic algorithm-ii. *International Journal of Geographical Information Science*, *25*(12), 1949–1969.

Caparros-Midwood, D., Barr, S., & Dawson, R. (2017). Spatial optimization of future urban development with regards to climate risk and sustainability objectives. *Risk Analysis*, *37*(11), 2164–2181.

Casalino, L. P., Devers, K. J., Lake, T. K., Reed, M., & Stoddard, J. J. (2003). Benefits of and barriers to large medical group practice in the united states. *Archives of Internal Medicine*, *163*(16), 1958–1964.

Church, R. L. (2002). Geographical information systems and location science [Journal Article]. *Computers and Operations Research*, *29*(6), 541 - 562. Retrieved from http://www.sciencedirect.com/science/article/pii/S0305054899001045

Crooks, V. A., & Schuurman, N. (2012). Interpreting the results of a modified gravity model: examining access to primary health care physicians in five canadian provinces and territories. *BMC Health Services Research*, *12*(1), 1–13.

Dahrouge, S., Hogg, W., Younger, J., Muggah, E., Russell, G., & Glazier, R. H. (2016). Primary care physician panel size and quality of care: a population-based study in ontario, canada. *The Annals of Family Medicine*, *14*(1), 26–33.

Daskin, M. S., & Dean, L. K. (2005). Location of health care facilities. In *Operations research and health care* (pp. 43–76). Springer.

Duckett, S., Breadon, P., & Ginnivan, L. (2013). *Access all areas: new solutions for gp shortages in rural australia* (Tech. Rep.).

Elkabalawy, M., & Al-Sakkaf, A. (2021). Decision support nsga-ii optimization method for resource-constrained schedule compression with allowed activity splitting. In *Journal of physics: Conference series* (Vol. 1900, p. 012016).

Exeter, D. J., Zhao, J., Crengle, S., Lee, A., & Browne, M. (2017). The new zealand indices of multiple deprivation (imd): A new suite of indicators for social and health research in aotearoa, new zealand. *PloS one*, *12*(8), e0181260.

Farahani, R. Z., SteadieSeifi, M., & Asgari, N. (2010). Multiple criteria facility location problems: A survey [Journal Article]. *Applied Mathematical Modelling*, *34*(7), 1689 - 1709. Retrieved from http://www.sciencedirect.com/science/article/pii/S0307904X09003242

Foley, R. (2002). Assessing the applicability of GIS in a health and social care setting: planning services for informal carers in east sussex, england [Journal Article]. *Soc Sci Med*, *55*(1), 79-96.

Foley, R., & Darby, N. (2002). Placing the practice: GP service location planning using GIS in Brighton & Hove. In *Geohealth 2002, victoria university of wellington, december 3-5th.*

Goodyear-Smith, F., & Janes, R. (2008). New Zealand rural primary health care workforce in

2005: more than just a doctor shortage. *Australian Journal of Rural Health*, *16*(1), 40–46.
Gordon, A., & Womersley, J. (1997). The use of mapping in public health and planning health services [Journal Article]. *J Public Health Med*, *19*(2), 139-47.
Graber-Naidich, A., Carter, M. W., & Verter, V. (2015). Primary care network development: the regulator's perspective [Journal Article]. *Journal of the Operational Research Society*, *66*(9), 1519-1532.
Griffin, P. M., Scherrer, C. R., & Swann, J. L. (2008). Optimization of community health center locations and service offerings with statistical need estimation [Journal Article]. *IIE Transactions*, *40*(9), 880-892.
Guagliardo, M. F. (2004). Spatial accessibility of primary care: concepts, methods and challenges [Journal Article]. *Int J Health Geogr*, *3*(1), 3.
Güneş, E. D., Melo, T., & Nickel, S. (2019). Location problems in healthcare. In *Location science* (pp. 657–686). Springer.
Güneş, E. D., Yaman, H., Çekyay, B., & Verter, V. (2014). Matching patient and physician preferences in designing a primary care facility network. *Journal of the Operational Research Society*, *65*(4), 483–496.
Hillsman, E. L. (1980). *Multiobjective location planning for primary medical services in rural iowa* (Tech. Rep.). Oak Ridge National Lab., TN (USA).
Hodgson, M. J., Laporte, G., & Semet, F. (1998). A covering tour model for planning mobile health care facilities in Suhum District, Ghana. *Journal of Regional Science*, *38*(4), 621–638.
Jaeggi, D. M., Parks, G. T., Kipouros, T., & Clarkson, P. J. (2008). The development of a multi-objective tabu search algorithm for continuous optimisation problems. *European Journal of Operational Research*, *185*(3), 1192–1212.
Kelly, E., & Stoye, G. (2014). *Does gp practice size matter? gp practice size and the quality of primary care* (No. R101). IFS Report.
Lapierre, S. D., Myrick, J. A., & Russell, G. (1999). The public health care planning problem: A case study using geographic information systems [Journal Article]. *Journal of Medical Systems*, *23*(5), 401-417. Retrieved from `https://doi.org/10.1023/A:1020585302393`
Leitch, S., Dovey, S. M., Samaranayaka, A., Reith, D. M., Wallis, K. A., Eggleton, K. S., ... others (2018). Characteristics of a stratified random sample of New Zealand general practices. *Journal of Primary Health Care*, *10*(2), 114–124.
Lopane, F. D., Barr, S., James, P., & Dawson, R. (2019). Optimization of resource storage location for managing flood emergencies [Conference Paper]. *ICONHIC 2019 Proceedings*. Retrieved from `https://iconhic.com/2019/2019/10/id-58-f-d-lopane-s-barr-p-james-r-dawson/`
Luo, W., & Qi, Y. (2009). An enhanced two-step floating catchment area (E2SFCA) method for measuring spatial accessibility to primary care physicians [Journal Article]. *Health Place*, *15*(4), 1100-7.
Luo, W., & Wang, F. (2003). Measures of spatial accessibility to health care in a GIS environment: synthesis and a case study in the Chicago region. *Environment and Planning B: Planning and Design*, *30*(6), 865–884.
Mazumdar, S., Bagheri, N., Konings, P., Chong, S., Jalaudin, B., Girosi, F., & McRae, I. (2019). Measuring relationships between doctor densities and patient visits: A dog's breakfast of small area health geographies [Journal Article]. *Applied Spatial Analysis and Policy*, *12*(3), 631-645. Retrieved from `https://doi.org/10.1007/s12061-018-9261-y`
Mazumdar, S., Butler, D., Bagheri, N., Konings, P., Girosi, F., Feng, X., & McRae, I. (2016). How useful are primary care service areas? Evaluating PCSAs as a tool for measuring primary care practitioner access [Journal Article]. *Applied Geography*, *72*, 47 - 54. Retrieved from `http://www.sciencedirect.com/science/article/pii/S0143622816300765`
Mazumdar, S., Feng, X., Konings, P., McRae, I., & Girosi, F. (2014). A brief report on primary care service area catchment geographies in New South Wales Australia [Journal Article]. *International Journal of Health Geographics*, *13*(1), 38. Retrieved from `https://doi.org/10.1186/1476-072X-13-38`

McGrail, M. R., & Humphreys, J. S. (2009). Measuring spatial accessibility to primary care in rural areas: Improving the effectiveness of the two-step floating catchment area method [Journal Article]. *Applied Geography*, *29*(4), 533 - 541. Retrieved from `http://www.sciencedirect.com/science/article/pii/S0143622809000034`

McGrail, M. R., & Humphreys, J. S. (2015). Spatial access disparities to primary health care in rural and remote australia [Journal Article]. *Geospat Health*, *10*(2), 358.

McLafferty, S. L. (2003). GIS and health care [Journal Article]. *Annual Review of Public Health*, *24*(1), 25-42.

Ministry of Health NZ. (2021a). (`https://www.health.govt.nz`, last accessed: 14.01.2021)

Ministry of Health NZ. (2021b). (`https://www.health.govt.nz/your-health/services-and-support/health-care-services/visiting-doctor-or-nurse`, last accessed: 14.01.2021)

Murad, A. (2004). Creating a GIS application for local health care planning in Saudi Arabia [Journal Article]. *International journal of environmental health research*, *14*, 185-99.

Panagiotis Mitropoulosa, I. G., Ioannis Mitropoulos. (2013). Combining dea with location analysis for the effective consolidation of services in the health sector [Journal Article]. *Computers and Operations Research*, *40*(9), 2241 - 2250. Retrieved from `http://www.sciencedirect.com/science/article/pii/S0305054812000202`

Parker, B. R., & Srinivasan, V. (1976). A consumer preference approach to the planning of rural primary health-care facilities. *Operations Research*, *24*(5), 991–1025.

Rahman, S.-u., & Smith, D. K. (2000). Use of location-allocation models in health service development planning in developing nations. *European Journal of Operational Research*, *123*(3), 437 - 452.

Raith, A., & Ehrgott, M. (2009, apr). A comparison of solution strategies for biobjective shortest path problems. *Computers and Operations Research*, *36*(4), 1299–1331.

Reuter-Oppermann, M., Nickel, S., & Steinhäuser, J. (2019). Operations research meets need related planning: Approaches for locating general practitioners' practices. *PloS one*, *14*(1), e0208003.

Reuter-Oppermann, M., Rockemann, D., & Steinhäuser, J. (2017). A GIS-based decision support system for locating primary care facilities. In S. Za, M. Drăgoicea, & M. Cavallari (Eds.), *Exploring services science* (pp. 210–222). Cham: Springer International Publishing.

Ribeiro, A., & Antunes, A. P. (2002). A GIS-based decision-support tool for public facility planning [Journal Article]. *Environment and Planning B: Planning and Design*, *29*(4), 553-569.

Rosero-Bixby, L. (2004). Spatial access to health care in Costa Rica and its equity: A GIS-based study [Journal Article]. *Social science and medicine*, *58*, 1271-84.

Schröder, L., Flägel, K., Goetz, K., & Steinhäuser, J. (2018). Mobility concepts and access to health care in a rural district in germany: a mixed methods approach [Journal Article]. *BMC Fam Pract*, *19*(1), 47.

Schuurman, N., Berube, M., & Crooks, V. A. (2010). Measuring potential spatial access to primary health care physicians using a modified gravity model. *The Canadian Geographer/Le Geographe Canadien*, *54*(1), 29–45.

Scott, K. M., Marwick, J. C., & Crampton, P. R. (2003). Utilization of general practitioner services in New Zealand and its relationship with income, ethnicity and government subsidy. *Health Services Management Research*, *16*(1), 45–55.

Shortt, N. K., Moore, A., Coombes, M., & Wymer, C. (2005). Defining regions for locality health care planning: a multidimensional approach [Journal Article]. *Soc Sci Med*, *60*(12), 2715-27.

Southern Health. (2021). (`https://www.southernhealth.nz`, last accessed 14.01.2021)

Standridge, C. R., Alan, A., Pritsker, B., & Delcher, H. (1978). Issues in the development of a model for planning health manpower [Journal Article]. *Simulation*, *31*(1), 9-13.

Steiner, M. T. A., Datta, D., Neto, P. J. S., Scarpin, C. T., & Figueira, J. R. (2015). Multi-objective optimization in partitioning the healthcare system of Parana State in Brazil. *Omega*, *52*, 53–64.

Tien, J. M., & El-Tell, K. (1984). A quasihierarchical location-allocation model for primary health care planning. *IEEE transactions on systems, man, and cybernetics*(3), 373–380.

Verter, V., & Lapierre, S. D. (2002). Location of preventive health care facilities. *Annals of Operations Research*, *110*(1-4), 123–132.

Walsh, S. J., Page, P. H., & Gesler, W. M. (1997). Normative models and healthcare planning: network-based simulations within a geographic information system environment [Journal Article]. *Health Serv Res*, *32*(2), 243-60.

Worldometer. (2021). (`https://www.worldometers.info/world-population/new-zealand-population/`. Last accessed: 14.01.2021)

Yanık, S., & Bozkaya, B. (2020). A review of districting problems in health care. In R. Z. Ríos-Mercado (Ed.), *Optimal districting and territory design* (pp. 31–55). Cham: Springer International Publishing.

Yin, H., Xiao, Y., Wen, G., & Fang, H. (2017). Design optimization of a new w-beam guardrail for enhanced highway safety performance. *Advances in Engineering Software*, *112*, 154-164.

Zhang, W., & Fujimura, S. (2010). Improved vector evaluated genetic algorithm with archive for solving multiobjective pps problem. *2010 International Conference on E-Product E-Service and E-Entertainment*, 1-4.

Zhang, Y., Berman, O., Marcotte, P., & Verter, V. (2010). A bilevel model for preventive healthcare facility network design with congestion. *IIE Transactions*, *42*(12), 865–880.

Zhang, Y., Berman, O., & Verter, V. (2009). Incorporating congestion in preventive healthcare facility network design. *European Journal of Operational Research*, *198*(3), 922–935.

Zhang, Y., Berman, O., & Verter, V. (2012). The impact of client choice on preventive healthcare facility network design. *OR spectrum*, *34*(2), 349–370.

Zhao, J., Ameratunga, S., Lee, A., Browne, M., & Exeter, D. J. (2019). Developing a new index of rurality for exploring variations in health outcomes in auckland and northland. *Social Indicators Research*, *144*(2), 955–980.