Decision support for the selection of optimal tower site locations for early-warning wildfire detection systems in South Africa

Andries Heyns[a,b,*], Warren du Plessis[c], Kevin M. Curtin[b], Michael Kosch[d,e,f,g], Gavin Hough[g]

[a]Humanitarian Logistics and Supply Chain Research Institute, Hanken School of Economics, Helsinki, Finland
[b]Laboratory for Location Science, Department of Geography, University of Alabama, USA
[c]University of Pretoria, Pretoria, South Africa
[d]South African National Space Agency, Hermanus, South Africa
[e]Department of Physics, Lancaster University, Lancaster, UK
[f]University of Western Cape, Bellville, South Africa
[g]EnviroVision Solutions, Durban, South Africa

* Corresponding author, contactable at andries.heyns@hanken.fi

Abstract

Effective early detection of forest fires can be achieved by specialised systems of tower-mounted cameras. Foresters and locals with intimate knowledge of the terrain traditionally plan the tower site locations – without the aid of computational optimisation tools. However, such knowledge and expertise may not be available to system planners when entering vast new territories. The process of selecting multiple tower sites from a large number of potential site locations with the aim of maximising system visibility of smoke above a prescribed region is a complex combinatorial optimisation problem. We present two recent applications of novel site-selection frameworks for tower-mounted *camera-based wildfire detection systems* (CWDS), which have been under development with guidance from experts from the South African-developed *ForestWatch* wildfire detection system. A novel single-site search framework determined alternatives for thirteen proposed sites in South Africa's Mpumalanga province, of which 6 alternatives were chosen over the initially proposed sites. The system-site selection framework was showcased in determining a four-camera CWDS layout in South Africa's Southern Cape – drastically improving on the detection capability of the layout initially proposed by technical experts.

Keywords: facility location, maximal cover, optimisation, wildfire

1 Introduction

Unexpected and uncontrolled wildfires spread rapidly and often turn into devastating natural disasters that affect the environment, ecosystems, economies and societies the world over. South Africa is no exception and suffers significant wildfire damage every year (Strydom and Savage, 2016). Wildfires are not only a threat to homes, families, and infrastructure, but also to the forestry assets of the South African timber industry. The South African forest sector employs roughly 165 900 workers and provides livelihood support to 652 000 of the country's rural population, and the government-run *Forestry Livelihoods Programme* is contributing to eradicating poverty (South African Government, 2009). A large percentage of South Africa's population is located in rural areas in these fire-prone forested regions and are especially vulnerable. It follows that the earliest possible detection of a wildfire is of critical importance. The sooner it is detected, the sooner suppressing action can be taken and the more manageable the size of the fire may be (Rego and Catry, 2006) – potentially minimising the loss of life, the scale of destruction, and the overall damage to the timber industry and affected livelihoods.

Early wildfire detection can effectively be achieved by *camera-based wildfire detection systems* (CWDSs) which comprise a number of specialised cameras that monitor the surrounding environment for smoke (Heyns et al., 2019; Martell, 2015). The cameras are mounted on top of towers that provide an elevated viewpoint above the terrain surface, resulting in improved visibility of the surrounding environment. Figure 1(a) shows a typical camera, while in Figure 1(b) a 32-m tower with a camera mounted on top is displayed. Human operators at dedicated workstations are alerted in order to validate a fire and, if validated, they notify fire protection agencies in order to initiate firefighting efforts (Heyns et al., 2019).

(a)

(b)

Figure 1 (a) Camera used in a CWDS; (b) a 32-m tower on top of which a camera is placed, with the solar power supply visible near the base of the tower *(Heyns et al., 2019)*.

The process of configuring the sites at which the towers of a CWDS are placed is critical to overall system detection potential with respect to the surrounding environment. Historically, site locations for CWDSs (or, more traditionally, watchtowers) have been planned without the use of computational optimisation tools by foresters and locals with intimate knowledge of the terrain. The areas that are considered to offer good candidate tower sites can be large and envelop expansive terrain surfaces (Heyns et al., 2019). The selection of a number of specific site locations – corresponding to the number of towers available – located on these large terrain surfaces poses a significant challenge. Simply identifying individual sites with good visibility cover of the surrounding environment would result in a number of cameras with good individual visibility. This is not a desired approach for *system* optimisation, where the overall detection potential depends on the *combined* visibility cover of *all* the cameras in the system instead of individual camera coverage.

When entering unfamiliar territories, the knowledge and expertise of foresters and locals may not be available to system planners. Selecting multiple tower sites to achieve comprehensive coverage becomes an even more daunting challenge in such instances. This can be alleviated by considering various combinatorial optimisation solution approaches which exist to identify multiple observation points with the aim of maximising *system* coverage – for examples, see (Bao et al., 2015), (Tong et al., 2009), (Zhang et al., 2019) and (Kim et al., 2004). The problem with relevant approaches from the literature, however, is that they are theoretical and do not address the *real-world* challenges associated with site selection problems. For example, these studies involve unrealistically small, square-shaped study areas with hypothetical test scenarios – and have no existing systems available in their study areas to at least provide some benchmark for their solution frameworks. Tower site selection approaches destined for large and more practically realistic areas exist (Eugenio et al., 2016), but are aimed at maximising single-site visibility cover with the potential for good system cover, rather than explicitly pursuing system coverage. A comprehensive framework aimed at the optimisation of *system* coverage achieved by CWDSs over *large* prescribed regions therefore remains absent from the literature.

To address the practical limitations from the literature, an optimisation framework for CWDSs has been under development with collaboration from *ForestWatch* (evsusa.biz) – a South African-developed CWDS with extensive operations in various critical regions in South Africa, and worldwide. The camera and tower in Figure 1 are, in fact, part of a ForestWatch CWDS in South Africa. Guided by their feedback and experience from an operational point of view and with the aim of maximising CWDS coverage, the intended purpose of this framework has evolved to a) determine multiple candidate CWDS tower-site layouts, b) within short timeframes (less than a week), and c) with minimal user input. Initial development of the framework investigated the wildfire-prone, mountainous region of Nelspruit in South Africa in which an existing twenty-tower CWDS served as a benchmark. The effectiveness of this benchmark CWDS has been proven by its daily detection numbers – in 2017 alone, the system logged 2786 alerts within the subscribed client area, and many more outside (Heyns et al., 2019). The smoke detection potential of the existing layout was compared to that which could be achieved by solutions determined by heuristic optimisation, and heuristic-obtained solutions outperformed the existing system (Heyns et al., 2019). Having such a successful existing system available as a benchmark for comparison together with guidance from technical experts from the region (who selected the sites for the existing system) allowed us to develop our approach with a level of detail and practical inspection which is missing from related studies in the literature. The study was expanded by investigation into the implementation of landform-based site selection (e.g. peaks, ridges, slopes) to improve our candidate site selection process (Heyns et al., 2020). The results allowed us to consider additional solution approaches which led to improved results within reduced computation times compared to our first attempts from Heyns et al. (2019). The problems above illustrated how using geographical information systems (GIS) together with our *multi-objective* (MO) optimisation approaches could drastically improve future CWDS system planning – not only in coverage maximisation but also in easing the actual decision-making processes.

The focus of this paper is to present two recent real-world tower site-selection problems that implement and build on the previous framework development (Heyns et al., 2019, 2020). First, a search for alternatives to thirteen separate towers proposed by ForestWatch technicians in the Mpumalanga province was investigated. This presented us with an opportunity to develop and implement a novel single-site selection framework for the identification of alternative sites for *individual* towers, as opposed to *system*-tower optimisation which had been the previous research focus. Single-site search approaches are not uncommon in the literature (Tabik et al., 2013; Zhang et al., 2019); however, approaches to search for *alternatives* for proposed sites do not exist. This also provided the first opportunity for practical implementation of landform-based site selection. The framework identified numerous alternatives for each of the thirteen sites proposed by ForestWatch, and six alternatives were eventually chosen above expert-selected sites. Furthermore, the single-site selection framework introduced new coverage evaluation criteria which had not been investigated in our previous work, nor in the related literature. These new coverage criteria are poised for consideration as additional objectives in future implementations of our framework.

Moving on from the single-site selection problem, the system-optimisation framework was showcased as fully-functioning and practical when it was applied to select sites for a new four-tower CWDS in South Africa's Southern Cape. Rural villages are found in this region and many of the local population are employed by the forestry sector. In 2017, in the town of Knysna (a mere 60 km away), one of South Africa's most devastating fires ever occurred (Forsythe et al., 2019). The study area exhibits similar vegetation and terrain to the Knysna area – a similar catastrophe occurring is thus a very real possibility and was one of the driving factors for the decision to install a CWDS here. Rapidly-determined layouts from our framework drastically outperformed the coverage achieved by sites initially proposed by technical experts with years of experience in forestry and tower site selection (and which required weeks of planning). One of our solutions was eventually selected instead of their initial layout, although two of its four tower sites were slightly altered by the decision-makers for practical purposes which are elucidated later. ForestWatch requested alternative layouts days before a contract proposal deadline – the fact that numerous superior and practically implementable layouts were obtained within such a short timeframe further substantiates why ForestWatch plans to implement the framework in future site-selection problems. Collaboration with decision-makers in determining tower sites before and after computational optimisation revealed interesting practical considerations and important guidelines for future work – a novel contribution to the literature related to similar problems in which the focus is overwhelmingly theoretical.

The paper opens with a summary of CWDSs in terms of their detection requirements and the factors that need to be considered in the planning of their tower site locations. The GIS component of our framework is elucidated, which includes a review of candidate site selection methods and how GIS is used in our framework for this purpose. Processes for the selection of final tower sites are then described, specifically related to optimisation methods and those implemented within our framework. The two problems presented in this paper are then discussed in terms of the project requirements and the data and methods used. The results are then presented, followed by a discussion and a brief conclusion.

2 Background

2.1 Camera-based wildfire detection systems

The type of clients that subscribe to CWDSs vary regionally. In South Africa they are forestry companies, while in the USA and Canada they are either local or federal government agencies. Subscription cost models vary between CWDS service providers. Some providers charge fixed fees per tower, while a provider such as ForestWatch calculates a subscription fee in relation to the total client-area coverage achieved. Coverage maximisation therefore not only contributes towards a CWDS's ability to detect smoke and initiate a response but may also result in increased subscription revenue. The client typically pays for the tower and equipment installation costs, providing further motivation to achieve comprehensive visibility coverage with the minimum number of required towers. Minimising the number of required towers to achieve optimal cover also results in reduced future expenses on maintenance and upgrades.

ForestWatch CWDSs detect smoke patterns and their effectiveness depends on their ability to observe smoke above the terrain surface (Heyns et al., 2019; Hough, 2007; Schroeder, 2005) – their algorithm is based on automated detection of aurora which they developed in Antarctica (Hough, 2007). This differs from the standard approach followed in surveillance system applications (including those related to CWDSs), where visibility is evaluated with respect to the terrain surface (Bao et al., 2015; Franklin, 2002; Kim et al., 2004; Tabik et al., 2013). In order to be visible from a camera, smoke needs to rise from the ground and typically needs to clear visibility obstruction from terrain and vegetation. Once smoke is identified, human operators are alerted and detection reports are sent. Detecting a smoke plume as low as possible above the terrain surface allows more rapid suppressing action to be taken after the onset of the fire. However, a camera's visibility of smoke is more likely to be obstructed by terrain and vegetation when the smoke is near the terrain surface or the fire is in a valley or behind a hill, as shown in Figure 2. A CWDS's overall detection potential therefore also depends on its ability to detect smoke at higher levels above the terrain surface (after clearing obstructions). Furthermore, CWDSs may be configured in such a manner that they achieve satisfactory visibility cover over *buffer zones* (Heyns et al., 2019). Buffer zones extend beyond the client boundaries since external fires may well encroach onto the client area and are also crucial to monitor.

Figure 2 Wildfire detected by the ForestWatch CWDS, displaying typical visibility obstruction caused by terrain.

2.2 Terrain and candidate site representation

CWDS site-selection optimisation requires an appropriate data environment within which to function. Raster data is employed extensively in the literature for solving facility location problems similar to the CWDS site-selection problem (Franklin, 2002; Heyns and Van Vuuren, 2018; Kim et al., 2004; Kwong et al., 2014) and represent the earth's surface and environmental information as uniformly spaced sample points across the terrain. A raster data representation of a hypothetical terrain surface is provided in Figure 3(a). The non-contiguous blue surface area is an example of suitable terrain identified for the placement of towers – subject to factors such as allowable geographical and administrative/municipal boundaries, and suitable terrain characteristics, or manual selection. The green surface area is an example of an area that requires monitoring and, in the context of this paper, is land belonging to forestry clients. The dots on the terrain surface are uniformly spaced satellite-sampled elevation data, which are used to generate the surface. The distance between neighbouring sample points is approximately 30 m at the highest resolution of raster data that is publicly available. The sites that may be considered for facility placement (the blue dots) collectively form the Placement Zone (PZ).

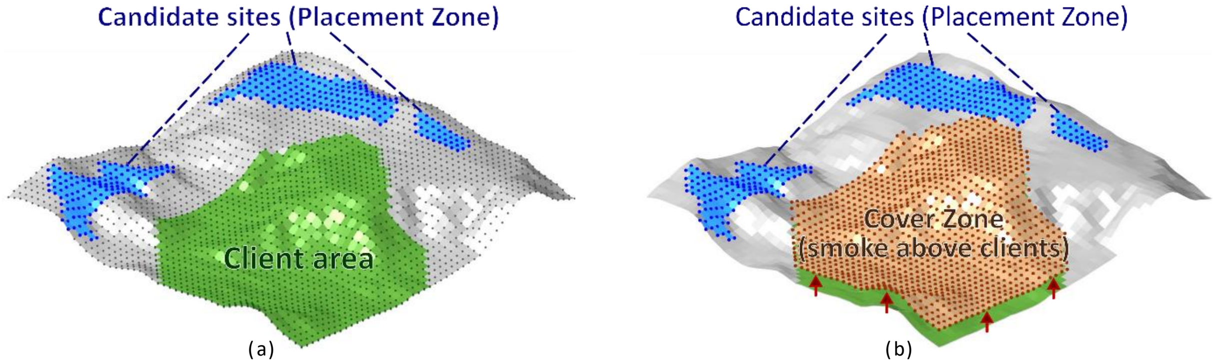

Figure 3 Raster data represent the earth's surface as uniformly spaced sample points *(Heyns et al., 2019)*. (a) Raster representation of a terrain surface with a PZ and client area; (b) raster representation of a Cover Zone above the client area.

2.3 Smoke layers and buffers

Detection potential is evaluated according to coverage achieved with respect to areas known as *Cover Zones* (CZs). In the context of CWDSs, a CZ is simply the rasterised terrain surface that falls within the client area (and extended to within some buffer boundary) raised to a specified height above the ground (simulating a layer of smoke) so that the system's potential for detecting smoke at that height may be evaluated. The buffer zone added to the smoke layer allows monitoring of the progress of fires outside the client area – these fires need to be monitored by ForestWatch, but client response is not necessarily required if their properties are not under immediate threat. An example of a CZ is illustrated in Figure 3(b) – the brown surface and markers above the client area.

In our framework, the CZs are considered at different heights above the terrain surface. One of the main added advantages of using more than one smoke height and buffer distance is that different tower-site combinations which contain sites at different locations are typically found to provide superior detection potential with respect to each CZ (Heyns et al., 2019). This leads to more diverse solutions and trade-off alternatives for decision-makers, i.e. more options (the benefits of this in the practical decision-making process are discussed in more detail later). Smaller buffer zones (0 to 500 m) are added to lower layers, intended for near-immediate detection and rapid response. The detection potential of higher layers gauges how well the system can detect smoke that has risen from the lower layers to clear obstructions to (potentially) be visible. Extended buffer zones (500 m to 4 km) are added to these higher layers. Figure 4 provides a visual description of the GIS processes involved in generating CZs using these methods.

The portion of a CZ that is visible from a camera is referred to as its *viewshed*, and is computed from a collection of line-of-sight queries between the camera and all the demand points within the CZ, limited by terrain interference and the camera's detection range (Kim et al., 2004; Nagy, 1994). A system viewshed of a CZ is then the merged viewsheds of all the individual cameras in the CWDS with respect to the CZ – i.e. the demand points in the CZ that are visible from at least one camera in the system.

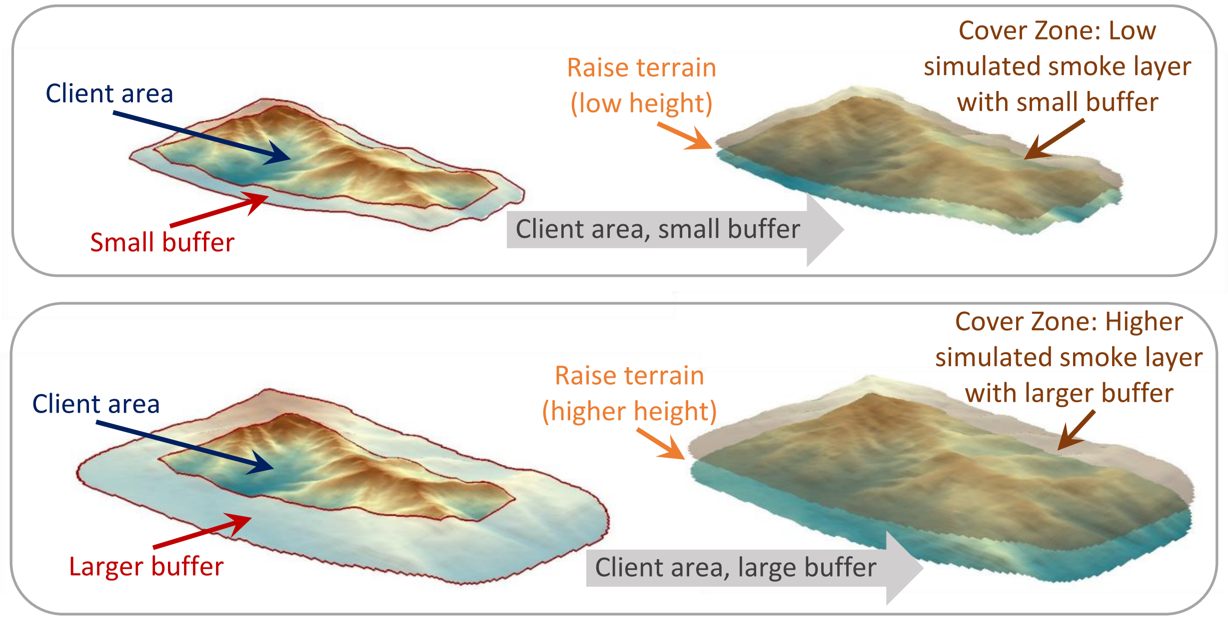

Figure 4 Buffers zones are added around the client area to monitor threatening external fires, and the combined client and buffer terrains are raised in order to simulate smoke layers at different heights above the terrain surface (the CZs). A small buffer with a low height is used to determine a CWDS's near-immediate detection capability (top image), while a larger buffer with a higher height is used to evaluate secondary detection potential (bottom image).

2.4 Candidate site identification

Identifying the candidate sites from which final tower sites are selected is a sensitive process. From an initial, "rough" pool of candidates, weaker sites may be identified and discarded, resulting in a stronger pool of candidates – while simultaneously reducing the computational complexity of the problem (Heyns et al., 2020). Caution should, however, be taken to avoid the possibility of removing good candidate sites by untested or excessive reduction techniques. Our approach requires the identification of candidate sites for single- and system-site searches, for which approaches have been described in the literature.

The least sophisticated candidate site selection approach is also the most arduous and time-consuming, namely manual candidate site selection. This approach is only relevant in small, hypothetical problem areas in which manual terrain inspection is a viable approach. Examples in the context of wildfire detection include the selection of 34 candidate sites by Zhang et al. (2019) and 30 candidate sites by Bao et al. (2015) – both study areas were smaller than 11 sq. km. and rectangular due to the theoretical focus of their work. The average ForestWatch system covers a surface area of well over 200 sq. km, and the practical implementation later in this paper has a client area of approximately 435 sq. km. These expansive terrains typically contain numerous mountains, hills and ridges, so manually identifying candidate sites would become impractical. This does not mean that manual site selection is impossible in such large areas – the existing towers in large regions monitored by ForestWatch have been selected manually. However, this has only been possible because technicians and experts with decades of experience in the regions were familiar with the terrain and because of historical lookouts and existing infrastructure in the areas already being well known. Even then, the manual site selection and inspection process took months (Heyns et al. 2019) and moving into unfamiliar terrain would pose an even greater challenge. Manual site selection is therefore not suitable for our purposes.

GIS approaches offer a relatively simple alternative and is suitable for identifying candidate sites in significantly large areas, while ensuring that the sites are practically feasible. Eugenio et al. (2016) searched for sites for manned watchtowers in a large area covering 46 000 sq. km in Brazil. GIS analyses were first used to identify land within feasible geographical and administrative/municipal boundaries, after which terrain feature classification analyses were used to identify ridges on mountains and hills. Areas on the terrain that were within suitable distances from roads were also identified. Sites that fell within a feasible terrain surface that satisfied all three criteria of feasible land, ridge features, and suitable road access resulted in a final set of candidates which were considered for watchtower placement. This method of site identification avoids the manual process followed by Bao et al. (2015) and Zhang et al. (2019) and is suitable for implementation in our approach. We use similar GIS approaches to those used by Eugenio et al. (2016) to identify suitable areas from which candidate sites may be selected (Heyns et al., 2019), discussed next. The disadvantage of the approach is that it results in unusually large numbers of candidate sites, but this is mitigated using combined heuristic approaches in our framework.

2.5 GIS for candidate site selection

The GIS component of our framework limits the terrain that lies within client boundaries to raster points which exhibit suitable characteristics for tower placement. First, terrain with a degree of slope over 12° (or 20%) should be avoided to ensure that tower installation may be performed without the need for excessive terrain alteration, in addition to ease of access on foot. Second, for transportation and general access purposes (*e.g.* construction and maintenance), a distance of 100 m or less to roads is deemed necessary. The criterion of proximity to power supplies has been considered, but solar power supplies are generally installed due to an inconsistent power supply system in the South Africa and theft (a solar power supply can be seen in Figure 1(b)). Nevertheless, access to power may yet be implemented in future problems – its (indirect) importance in practice was illustrated in the decision-making process of experts related to the problems presented in this paper. Software such as the commercially available ArcGIS 10.5.1 can be used for the purpose of terrain and site analysis. Feasible slope sites are determined with 30 m resolution raster elevation data and the ArcGIS slope tool, while road-accessible sites are determined with roads data obtained from the clients in the study area and the ArcGIS Euclidean distance analysis tool. The resulting PZ consists of sites which satisfy both these requirements. These criteria and analyses were shown to be realistic when it was found that the sites of 26 towers in the benchmark Nelspruit CWDS were all located at sites which satisfied these requirements (Heyns et al., 2019).

Reducing the size of the PZ to landforms that are typically associated with superior visibility – more generally referred to as the *reduced observer strategy* (Rana, 2003) – is also integrated in our framework (Heyns et al., 2020). In the related literature, ridges and peaks are consistently considered to offer superior observer visibility compared to sites classified otherwise (Franklin and Clark, 1994; Kim et al., 2004; Lee, 1994; Rana, 2003). Reducing the PZ to such landform types reduces the number of candidate sites and results in reduced combinatorial complexity, while it has also been shown that the approach results in improved solution quality because superior sites are considered in the search process and inferior ones are avoided. Our framework implements *geomorphons* – these are pre-defined terrain patterns that are matched to land surfaces according to similarities in their geometry (Jasiewicz and Stepinski, 2013). In a single, simple execution (requiring a single line of code) the geomorphon tool can identify ten significant landform classes: flats, peaks, ridges, shoulders, spurs, slopes, pits, valleys, footslopes and hollows, as illustrated in Figure 5. The geomorphon classification approach is implemented in our framework due to its simplicity and availability in open-source software, and its proven practicality in a variety of recent problems (Di Stefano and Mayer, 2018; Djurdjevac Conrad et al., 2018; Harmon et al., 2018; Luo and Liu, 2018). All geomorphon classifications in this article were processed in the GRASS 7.4.0 software environment. An example of the results of a geomorphon classification is provided in Figure 6 for the terrain surrounding a tower site in (a), with the corresponding geomorphons in (b).

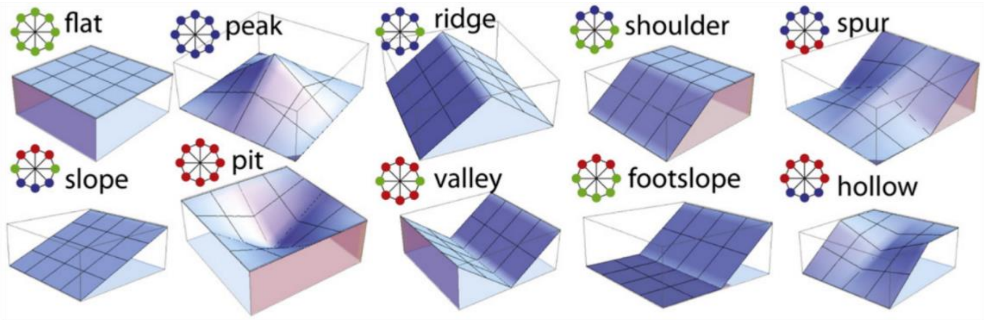

Figure 5 Terrain landform classifications of (Jasiewicz and Stepinski, 2013). The colours of the patterns alongside each class indicate differences in elevation with respect to the centre point – green indicates same height, red indicates higher, blue indicates lower.

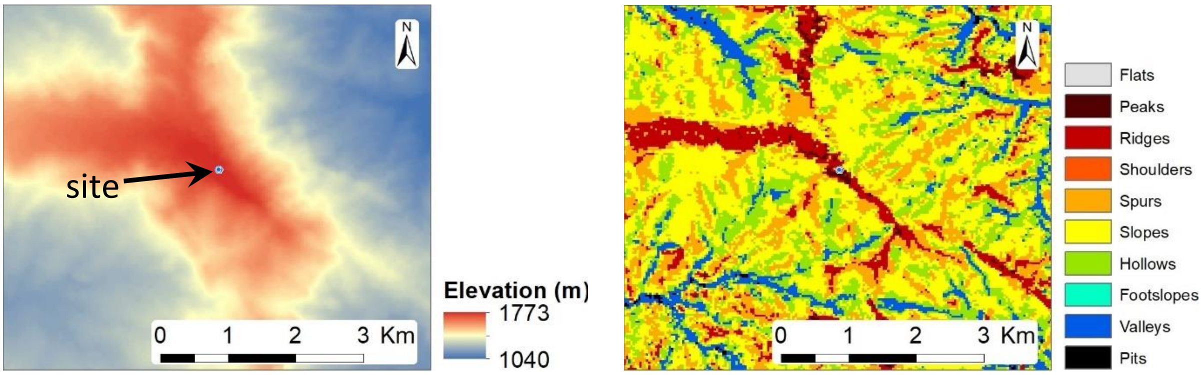

(a) (b)

Figure 6 (a) Terrain elevation around a proposed site location, and (b) the corresponding geomorphon landform classification of the surrounding terrain.

The strategy of selecting candidate sites according to landforms such as peaks and ridges should, however, be approached with caution because a level of uncertainty is introduced which may result in good sites being discarded (Romero and Clarke, 2018). Therefore, to avoid the unsubstantiated implementation of geopmorphons it was decided to first analyse the terrain feature classes at 165 ForestWatch towers from systems in Mpumalanga Province in South Africa (93 towers), Douglas County in the state of Oregon, USA (31 towers) and the central region of Saskatchewan Province in Canada (41 towers). This was the first time that such a *practical* site classification exercise has been performed for *existing* facilities, as opposed to some traditional analyses in a theoretical context related to terrain only – e.g. those performed by Kim et al. (2004) and Rana (2003). It was found that 136 (or 82%) of the towers were sited at peak or ridge sites as classified by the geomorphon approach, while those that are sited otherwise are never far away from peaks or ridges (less than 175 m) (Heyns et al., 2020). Discussions with ForestWatch technicians revealed that some towers are located at sites classified other than peaks or ridges because even though the peak or ridge sites would actually be preferred, they are sacrificed for nearby alternatives due to factors such as ground condition and accessibility. Nevertheless, peaks and ridges are the go-to sites according to geomorphon landform analysis and according to technicians. Identifying candidate sites that are limited to peaks and ridges should therefore provide decision makers with sites that are either a) practical and selected for final implementation, or b) sufficiently close to nearby alternatives which may be considered more suitable for practical reasons.

2.6 Final tower site selection

2.6.1 Practical processes

In practice, the selection of final tower sites is an iterative process between CWDS providers and clients and/or detection agencies involved in protecting a specific region, and will often aim to finalise and deploy a CWDS layout in time for a fire season. In the past, this lengthy process (without significant use of any computerised support) has led to suitable strategies not being agreed upon in time for a CWDS's deployment prior to a fire season, resulting in the deployment being delayed another year and the vulnerable region being left to endure one more season with existing, outdated and inferior detection ability – or none at all. When determining a suitable CWDS layout, the client should be pleased with the predicted detection coverage, while the cost of the installation and operation of towers is also important. The towers are generally sited on the client's property and the client may dislike one or more sites for subjective reasons. Such scenarios require technicians to conduct the site-selection process anew, with the possible requirement of finding alternatives for *all* the proposed sites (i.e. no sites retained). This is because moving a single undesirable site not only influences its individual coverage, but that of the system – resulting in its relocation requiring the relocation of another tower to compensate for the changes in system coverage, followed by further relocations to offset the second relocation's effects, and so on. Naturally, reducing the number of towers required to achieve satisfactory cover is preferred, to reduce installation and operational costs, and to reduce the physical client property required for installation.

Due to factors such as those mentioned above, decision-makers would benefit from obtaining *multiple* CWDS layout alternatives from a decision-support framework. They may browse through proposed coverage maps achieved by these layouts and identify those which they consider to offer the best client coverage. The tower locations of their preferred layouts may then be investigated – this may be performed virtually with tools like Google Earth to perform a basic assessment, followed by physical site inspections if necessary. In the event that the client dislikes one or more sites in a preferred layout, alternatives which do not include the undesirable sites may be investigated. Furthermore, if specific sites that are considered undesirable by the client appear in multiple preferred layout proposals, it should be possible to remove them from consideration and *rapidly* repeat the optimisation process anew, providing new coverage maps achieved with more suitable CWDS layouts. Our framework has been developed with these key requirements in mind.

After the final tower sites have been agreed upon, a suitable tower height at each site needs to be determined. Extensions to base tower heights are normally added because an increase in tower height improves overall smoke detection potential by allowing a camera to see over obstructions. Base structures typically stand 12 m tall in South African projects and height increases are achieved by adding one or more extensions to these, normally in 3 m increments (Heyns et al. 2019). An increase in tower height at a site depends on a) whether an increase in tower height is required for the camera to rise above the canopy of surrounding trees, b) the actual need for an increase in height from the base, depending on client coverage already achieved from the base height, and c) whether the demands of an increase in structure size and support (in terms of the tower foundation and stabilisation wires that increase in span as tower height increases) can be accommodated at the site. For the remainder of this paper, it is assumed that site searches are performed with 12-m towers only and the focus is on site-selection only. Furthermore, a camera range of 8 km is generally used by ForestWatch for site search analyses in South Africa and will also be used for site search analyses in this paper (the cameras have a visible range of well beyond 8 km, but 8 km is used for contractual purposes).

2.6.2 Computational methods

Theoretical research into the evaluation of multiple candidate viewpoints' viewsheds from which a superior site may be identified is available in the literature (Lee, 1994; O'Sullivan and Turner, 2001; Tabik et al., 2013). Such computational approaches provide a platform for *single*-site searches. Zhang et al. (2019) perform sequential single-site searches after determining their 36 candidate sites for wildfire detection purposes. First, the viewsheds and covering percentages of each candidate site is determined and the single site with the best coverage is selected for tower placement. The selected site is removed from the set of candidates, the demand region is updated by removal of the demand area covered by the new tower, and then the next tower site is determined by finding the next candidate site with the best coverage over the updated demand. This process is then repeated until a budget limit has been reached, or until acceptable cover has been determined. While this final site-selection process is more user-friendly than a manual approach, it is a repeated single-site search destined for incremental expansions to existing towers – the process is not aimed at system-optimisation. This approach is therefore not considered suitable for our requirements, but an approach similar to their sequential one was implemented in our single-site alternative searches. However, compared to the literature in which a superior site is sought based on the coverage determined with respect to the surrounding terrain surface, our requirements are unique. Instead, we search for the best *alternatives* within a local proximity to an already proposed site – a requirement not previously considered in the literature. Furthermore, we evaluate alternatives according to three covering objectives in our single-site alternative search – resulting in more than one alternative – whereas the related literature focusses on identifying the single point with the best visibility according to a single covering criterion.

Eugenio et al. (2016) followed an interesting approach which essentially combines multiple localised single-site searches for overall system-optimisation. Their 46 000 sq. km study area was sub-divided into uniform square cells with the sides measuring 15 km, 17.5 km, and 20 km in separate analyses. The site with the highest altitude within each cell was selected as a watchtower site (constrained to suitable land boundaries, ridge features, and road access). In this manner, they were able to rapidly determine over 130 towers at a time in separate analyses, with the assumption that each tower in each cell would provide a good contribution to the overall system coverage of all the towers combined. The disadvantage of the approach is that the final selection of sites is limited to a single one within each cell and is ultimately based upon altitude. As is the case with manual site selection, merely identifying multiple sites which may provide individual watchtowers with good visibility does not guarantee good overall system cover. Their approach does not consider overall system coverage in the site selection process, and system coverage is only determined post-site selection. The approach may also result that superior system-sites are discarded within a cell because of it not providing the best perceived individual cover in the cell. Furthermore, high altitude alone does not necessarily ensure good visibility from a site and its relationship with its surrounding environment, and towers, is equally important (Franklin and Clark, 1994; Misthos et al., 2018) – especially when the aim is system coverage optimisation. This process is therefore not considered for our framework.

Bao et al. (2015) investigated the use of integer-linear programming and a genetic algorithm for "true" *system*-optimisation. They obtained candidate CWDS layouts comprising between six and sixteen watchtowers in single runs selected from thirty candidate sites – determined specifically with respect to a system coverage maximisation objective. While their theoretical study was conducted in an impractically small area of 10 sq. km and their sites were manually selected, the computational approaches that they followed are perfectly suited to our framework. Our problems, however, require additional heuristic approaches as a result of our significantly large, real-world territories and the resulting computational complexity. These approaches are discussed next.

2.7 Optimisation of tower site selection

2.7.1 Pareto-optimal solutions

Our objectives are to maximise the percentage of points in each CZ which are visible, i.e. to maximise visibility with respect to different smoke layers. Candidate CWDS layouts are evaluated by objective functions which calculate their detection potential with respect to each CZ. This translates to a single point in *objective function space* for each candidate layout (i.e. candidate solution), as is illustrated in Figure 7 in which a number of candidate layouts have been evaluated. The example in Figure 7 considers two CZs, which correspond to the two objectives on the axes, but the same principles apply for three or more objectives.

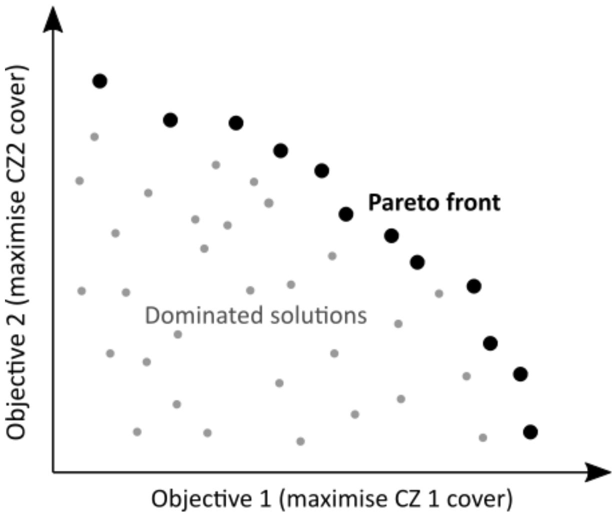

Figure 7 The notions of solution domination and of a Pareto front in objective function space.

In MO optimisation, solutions such as those in the figure are classified either as *non-dominated* (superior) or *dominated* (inferior) solutions (Zitzler et al., 2004). Dominated solutions are avoided, since for each dominated solution there exists at least one non-dominated solution that is equally good with respect to all the objectives, and is better in at least one. Amongst the solutions in the non-dominated set, each solution outperforms another in at least one of the objectives while simultaneously being weaker in at least one of the others. The set of non-dominated solutions exhibit superior trade-off alternatives to the dominated solutions, and form what is commonly known as the *Pareto-optimal front*, or simply the *Pareto front*, as may be observed in Figure 7 (Zitzler et al., 2004). Only the solutions on the Pareto front need to be presented to decision makers because of their superior quality.

2.7.2 Stage 1 - heuristics

The set of all possible solutions to a problem, *i.e.* all the possible candidate CWDS layouts on the terrain, is called the *solution space*. If N_t and N_s denote the number of towers available for placement and the number of feasible sites, respectively, the number of possible solutions, N_p, is

$$N_p = \binom{N_s}{N_t} = \frac{N_s!}{N_t!\,(N_s - N_t)!} \qquad (1)$$

The number N_p is imposingly large in problems such as those investigated in this paper, because the large scale of territories in which ForestWatch operate and the choice of a GIS-based candidate site identification approach. The pursuit of the *exact* (true) Pareto front in such instances (and in other MO facility location problems that include covering objectives) involves a significant computational challenge and become prohibitively large to solve within realistic computation times (Jia et al., 2007; Owen and Daskin, 1998; ReVelle and Eiselt, 2005; Tong et al., 2009; Xiao et al., 2002). Furthermore, not only is the search combinatorially complex in terms of the number of candidate sites and towers to place, but the computation of viewsheds (and system-viewsheds) in visibility-cover location problems imposes an additional and time-consuming computational burden (Heyns and van Vuuren, 2016).

In instances such as these, powerful heuristics are often employed in order to *approximate* the set of solutions on the Pareto front within realistic computation times (Bao et al., 2015, Xiao et al., 2002; Zitzler et al., 2004). These heuristics explore promising regions of the solution space in order to determine solutions that are approximately Pareto-optimal, and in the process avoids the computationally expensive consideration of solutions in inferior regions of the solution space. It has been demonstrated in the literature that heuristics are well capable of determining the true Pareto front (Heyns and Van Vuuren, 2018; Heyns and van Vuuren, 2016; Kim et al., 2008). *Multi-Objective Evolutionary Algorithms* (MOEAs) are able to approximate a diverse set of trade-off solutions on the Pareto front in a single run (Fonseca and Fleming, 1993; Purshouse and Fleming, 2003) and are also known to achieve good results fast (Alp et al., 2003). Examples of the application of MOEAs to solve problems similar to CWDS planning include the placement of transmitters (Meunier et al., 2000; Raisanen and Whitaker, 2005), wind turbines (Kwong et al., 2014; Yamani Douzi Sorkhabi et al., 2016), and observation equipment (Bao et al., 2015; Heyns and Van Vuuren, 2015; Kim et al., 2004; Tong et al., 2009). However, when a problem is sufficiently large and the location of the true Pareto front is unknown, there is no guarantee that the obtained solutions are on or even near to the true Pareto front.

The *non-dominated sorting genetic algorithm-II* (NSGA-II) is a popular MOEA that is classified as a *genetic algorithm*, in which a candidate CWDS layout is represented as a *chromosome* string of N_t feasible tower site numbers (Deb et al., 2002; Heyns et al., 2019). Site numbers are indexed for all the sites within the PZ's raster representation – typically derived with respect to row and column indices (Heyns and van Vuuren, 2016). For example, a chromosome [22, 115, 698, 739] represents a candidate CWDS with four towers located at sites 22, 115, 698 and 739. Evolution-inspired selection processes and modification operators are iteratively performed on a randomly generated population of such candidate CWDS chromosomes. The process is repeated until some termination criterion is met (Deb et al., 2002). One typical termination criterion is when successive populations fail to significantly improve on the solution quality of previous generations (Heyns, 2016). More detailed descriptions of the NSGA-II as used for our purposes are available in the literature (Heyns et al., 2019; Heyns and van Vuuren, 2016).

The large scale of territories in which ForestWatch operate and the implementation of a GIS-based candidate site selection approach instead of a manual one leads to a large number of candidate sites to consider in our problems – especially when compared to similar site optimisation problems in the literature (Bao et al., 2015; Kim et al., 2008, 2004; Tanergüçlü et al., 2012). The addition of viewshed-based covering objectives further adds to this computational complexity. Unique to our framework is the implementation of our *multi-resolution approach* (MRA) (Heyns, 2016; Heyns and van Vuuren, 2016) which alleviates this computational burden. The MRA is a recent, novel optimisation tool that was specifically developed for geospatial facility location problems with unusually large solution spaces such as those faced by ForestWatch. It has been shown that implementation of the MRA results in little or no reduction in solution quality, and in some instances can even lead to improved solution quality within drastically reduced computation times (Heyns, 2016; Heyns and van Vuuren, 2016).

The MRA simplifies the site search by first determining candidate layouts using a low-resolution grid of the candidate sites extracted from the high-resolution ones included in the feasible PZ area – effectively reducing the number of candidate sites. The NSGA-II is then run to approximate the Pareto front using this low-resolution PZ. The sites that are included in the solutions from this low-resolution Pareto front approximation are considered to be indicative of local regions which contribute favourably to optimal system coverage and merit further exploration. Thus, a finer resolution is used to intensify the search in the regions around these sites with additional optimisation runs. This is achieved by taking the low-resolution Pareto-front sites together with their high-resolution local neighbours, and pooling them together into a high-resolution pool of candidate sites – i.e. a new PZ.

Two resolutions have been used in our framework development (Heyns et al., 2019, 2020) and the real-world CWDS optimisation problem presented later in this paper. The first, lower resolution uses a spacing of approximately 90 m between neighbouring sites in the PZ (from the initial, higher resolution 30-m spacings). Then, around all the sites in the low-resolution Pareto front approximation, the feasible sites within a 5×5 raster-point neighbourhood at spacings of 30 m are selected and included in the high-resolution PZ. The algorithm is then run again with consideration of this high-resolution PZ. As an example of the initial reduction in computational complexity that may be achieved, the number of feasible sites in the PZ from our first study was reduced from 741 813 at 30-m spacings down to 82 547 at 90-m spacings (Heyns et al. 2019). The MRA also reduces the number of required viewshed computations and their associated computation time requirements, because the search is limited to promising regions and weaker ones are avoided (Heyns and van Vuuren, 2016). Most importantly, the MRA identifies and focuses on regions which contain sites that contribute to good overall *system* cover and not just on individual sites with good visibility.

2.7.3 Stage 2 - additional optimisation

The first optimisation stage as described above entails using the MRA-NSGA-II at two resolutions to determine multiple CWDS layouts. The final candidate layouts include multiple strong sites, but the unusually large size of our solution space and the approximation characteristics of the heuristic approach still do not guarantee that the solutions are optimal or even near-optimal. We therefore perform an additional optimisation stage.

The additional optimisation stage does not focus on determining new, alternative sites to those obtained in the first stage, but instead focusses on searching for improved site-combinations of these sites. These strong sites included in the Pareto-front approximations from the first stage are pooled together into a new PZ. This relatively small post-heuristic PZ then serves as input into the second stage's optimisation process in which additional runs can be performed – using either heuristics or ILP – without implementation of the MRA. Using the NSGA-II during this additional stage has been shown to result in significant improvement in the solution quality of CWDSs (Heyns et al., 2019). The number of sites in the new PZ is typically such a comparatively small number (compared to the size of the original PZs) that it has also introduced the possibility to implement an ILP weighted-sum approach as an alternative to heuristics in the second stage (Heyns et al., 2020). An ILP approach had not been considered previously, because ILP solver software are sensitive to the size and complexity of the problems which they can solve – heuristics, on the other hand, can attempt to find solutions to any size problem.

Commercial ILP software packages (e.g. CPLEX and Gurobi) take as input an ILP formulation of an objective function and constraints and return a single solution. Once the problem becomes multi-objective, the objectives are often weighted and summed together into a single objective function in order to satisfy the single-objective limitation of these software packages (Cohon, 1978; Murray et al., 2007). An approximation to the Pareto front is traced out by varying objective weights in multiple runs. This method provides a straightforward approach to solving MO optimisation problems because the ILP formulations are relatively simple to provide as input when compared to the requirements of heuristics such as the NSGA-II – which include sophisticated code and multiple parameters that require iterative tuning and an intimate knowledge of their effects. A strong characteristic of the weighted-sum approach is that the end-points of the Pareto front, which optimise with respect to only a single objective while ignoring the others, may be determined exactly. These end-points provide an indication of where the true Pareto front lies – and avoids a known weakness of MO heuristics which struggle to reach end-point regions (Kim et al., 2008).

The weighted-sum approach does, nevertheless, hold disadvantages. Evenly distributed weights may result in an unevenly distributed Pareto front approximation, and while truly optimal solutions can be found for the specific weight combinations, there is no guarantee that the returned solutions are on the true Pareto front (Khan and Rehman, 2013; Marler and Arora, 2010). Furthermore, assigning suitable weights to the objectives is a laborious and sensitive iterative process, and multiple runs are required in order to approximate the Pareto front (Marler and Arora, 2004; ReVelle and Eiselt, 2005). The weighted-sum approach remains a popular choice to solve MO optimisation problems from various applications (Machairas et al., 2014; Xia et al., 2019; Yao et al., 2018). The major advantage of this approach, in a practical sense, is that the user is able to specify the desired number of solutions in the form of the number of weight combinations. Heuristics often generate an impractically large number of solutions on the Pareto front approximation, which are unrealistic to present to decision-makers (Heyns et al., 2019) and require further analysis to be reduced to a manageable number (Heyns, 2016; Mavrotas, 2009).

The ILP formulation of the problem is now presented and is based on the *Maximal Covering Location Problem* (MCLP), first proposed and formulated by (Church and ReVelle, 1974). The CWDS planning problem includes multiple covering objectives evaluated with respect to multiple CZs, for which a multi-

CZ formulation of the MCLP was introduced by Heyns et al. (2020). The parameters used are listed below.

N_t	denotes the number of towers available for placement.
N_c	denotes the number of CZs.
s	denotes the number of candidate sites in the PZ.
d_c	denotes the index of demand points in CZ c, where $c \in \{1, \ldots, N_c\}$.
N_{d_c}	denotes the number of demand points in CZ c.
$\mathbb{N}_{d_c}$	denotes the subset of sites in the PZ from which demand point d_c in CZ c is visible.
x_s	is 1 if a tower is placed at site s, and 0 otherwise.
y_{d_c}	is 1 if a demand point d_c is covered, and 0 otherwise.

The objective is to:

$$\text{maximise} \quad V_c = \sum_{d_c} y_{d_c} \quad \forall \; c \in \{1, \ldots, N_c\} \tag{2}$$

subject to the constraints:

$$y_{d_c} \leq \sum_{s \in \mathbb{N}_{d_c}} x_s \tag{3}$$

$$\sum_{s} x_s = N_t \tag{4}$$

$$x_s \in \{0,1\} \tag{5}$$

$$y_{d_c} \in \{0,1\} \tag{6}$$

The objective in (2) is to maximise cover with respect to each CZ $c \in \{1, \ldots, N_c\}$. The constraint in (3) allows a demand point d_c to be covered ($y_{d_c} = 1$) only if one or more cameras are placed at sites in the set $\mathbb{N}_{d_c}$. Constraint (4) ensures that exactly N_t towers are placed, while constraints (5) to (6) specify binary requirements on the auxiliary variables.

To arrive at the weighted objective function, the N_c objectives in (2) can be reduced to a single function using a weight, w_c, for each CZ. The objective is then to

$$\text{maximise} \quad V = \sum_{c} w_c \frac{100}{N_{d_c}} \sum_{d_c} y_{d_c} \tag{7}$$

The objective in (7) is subject to the same constraints (3) to (6), enforced with respect to all CZs. The fraction is included in the objective function to reflect the maximisation of the *percentage* of cover achieved with respect to each CZ, so that the objective function is not biased towards larger CZs with more demand points.

3 Data and methods

3.1 Single-site selection problem

3.1.1 Problem description

ForestWatch requested assistance in the selection of a number of sites in the Mpumalanga Province in December 2018. The problem did not require system optimisation – ForestWatch provided thirteen sites proposed by planners for which individual alternatives were sought. This served as an evaluation of the proposed sites, and to provide ForestWatch with alternatives if there were any that were significantly better than the proposed ones. This provided a practical opportunity to refine the GIS component of our framework by implementing the exploitation of landforms for the selection of superior sites, as earmarked for later use in system optimisation problems. The locations of the proposed sites are shown by red markers in Figure 8, along with the existing towers in the region (orange markers), the client areas already covered by the existing towers (the green surface area) and the client areas that are not covered (the black surface areas). The cover achieved by the existing towers outside of the client areas is indicated in blue in the figure (using an 8 km detection range).

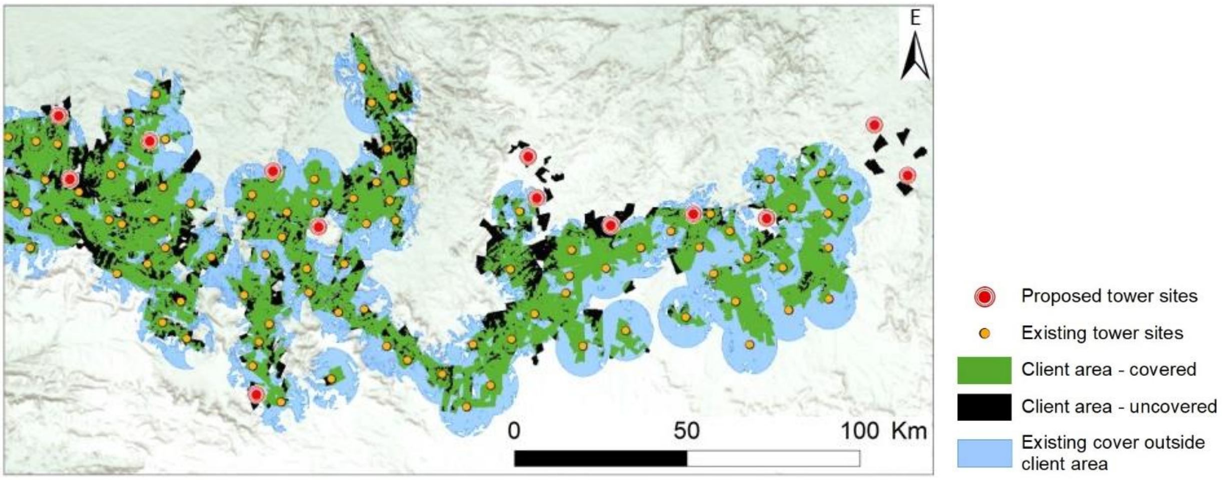

Figure 8 Thirteen sites (the red markers) that were proposed by ForestWatch, with the aim of providing additional cover over client plantations and surrounding areas.

3.1.2 Single-site solution framework

As previously discussed, searching for superior sites within a given region is not uncommon, but searching for alternatives for a proposed site was a challenge for which we required to develop our own framework. Sites that were classified as peaks or ridges within a 2 km radius around each of the proposed towers were identified as candidate alternative sites. A 2 km radius was agreed upon with decision makers as this is typically the maximum extent around which they would consider and search for alternatives in real-world site searches (4 km was used for one of the thirteen sites because of site-specific requirements outlined by technicians). Road accessibility was not considered here, because obtaining the roads layers from multiple clients in the region (with data that are typically conflicting between clients and/or outdated) was a task that was too laborious to complete and verify within the short timeframe that was available. Furthermore, it was decided to omit the consideration of suitable degree of slope that was considered in previous work (Heyns et al., 2019), because peak and ridge sites inherently exhibit low degrees of slope, as observed in Figure 5 (this was also assumed for the system-site optimisation problem presented later).

Exhaustive searches were performed using the identified candidate sites around each proposed tower, with the goal of providing multiple alternatives. To ensure this, three covering criteria were used to evaluate each candidate site, namely a) total cover achieved (client and outside), b) total client cover achieved (within client boundaries only), and c) total new client cover achieved (existing blind spots in the client area). Since the aim here was single-site optimisation, no layout alternatives were required, so only one smoke layer height of 30 m was considered for site evaluation. Recall that multiple smoke layers serve the important purpose of returning layouts with different tower-site location combinations, but this is not required here and the three covering objectives (at 30-m heights) are considered sufficient to return alternatives. Our framework is described next, and is illustrated in Figure 9 – repeated with respect to each covering criterion. As seen in the figure, once all candidate alternatives around each proposed site is identified and evaluated with respect to a criterion, two alternatives are identified. This is achieved by identifying the best-performing alternative with respect to a criterion – alternative 1 – after which this site and all others within 500 m are removed from the candidates, and the second-best alternative is identified from the remaining alternatives – alternative 2. The requirement of at least 500 m between the first and second-best alternatives is enforced to ensure that neighbouring sites are not proposed as first and second-best alternatives (neighbouring sites typically exhibit similar visibility results). This also ensures diversity in the locations of alternative site locations – and therefore more alternatives for decision-makers.

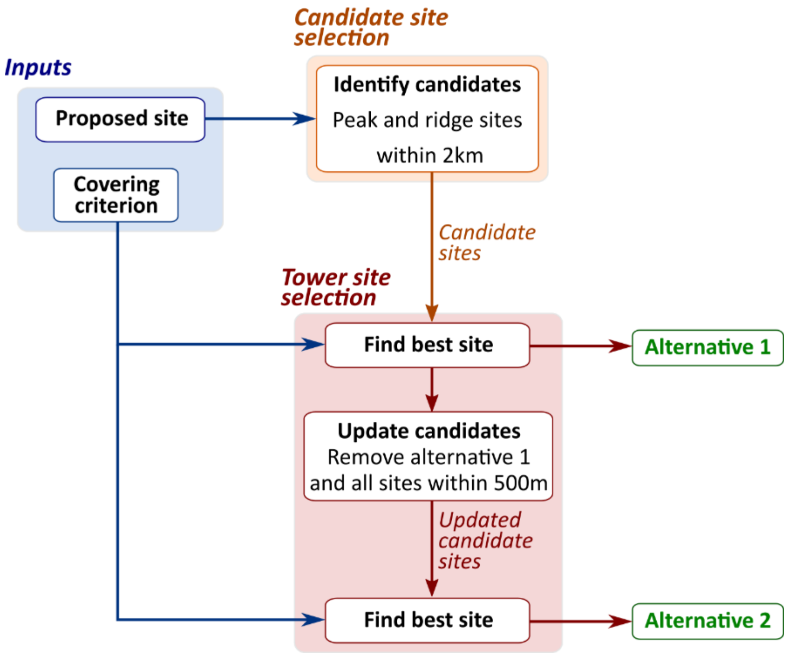

Figure 9 The single-site search framework used to determine alternative sites around a proposed site, determined with respect to a covering criterion – the framework is repeated for each criterion. In our application, three criteria were considered, and two alternatives were sought with respect to each criterion.

3.2 System-site selection problem

3.2.1 Problem description

In May 2019, ForestWatch requested optimal CWDS site locations to provide coverage to a forestry client in South Africa's Southern Cape (a total of 435 sq. km of client property), 60 km away from where the devastating 2017 Knysna fires occurred (Forsythe et al., 2019). Alternatives were sought to compare to a four-site layout that had been determined by ForestWatch technicians following weeks of speculation and physical site exploration. They had to propose a layout to their client within less than a week and requested an evaluation of their proposed layout and an additional investigation to identify possible superior alternatives. This meant that there was only enough time to implement the first stage of the optimisation framework (MRA-NSGA-II) to obtain alternatives, without any additional optimisation runs as has been performed in previous work (Heyns et al., 2019, 2020). As will be shown later, this did not have any significant impact on the solution quality of the layouts.

3.2.2 Preliminary analyses

The client boundaries are displayed in **Error! Reference source not found.**(a), in addition to the terrain surface that lies within 100 m from roads, indicated in grey. Roads data were obtained for those roads from which it was permissible to place towers – although some roads lie outside and between the two large client areas, permission was granted to consider this area for site placement due to agreements with ForestWatch and local authorities. Geomorphons were determined for the terrain surface, illustrated in Figure 10(b) (note that the legend only shows the main visible landform types and the others are not shown because of their scarcity). All sites that were identified to be within 100 m from roads and classified as peaks or ridges by the geomorphon approach were included in the final PZ, illustrated in

Figure 10(c). This PZ contained 46 483 sites (30-m resolution), which was reduced down to 5 172 at the lower 90-m resolution for the MRA.

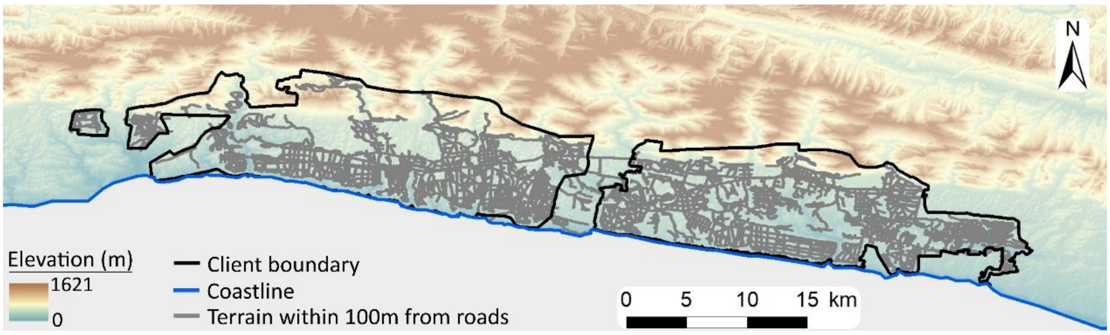

(a)

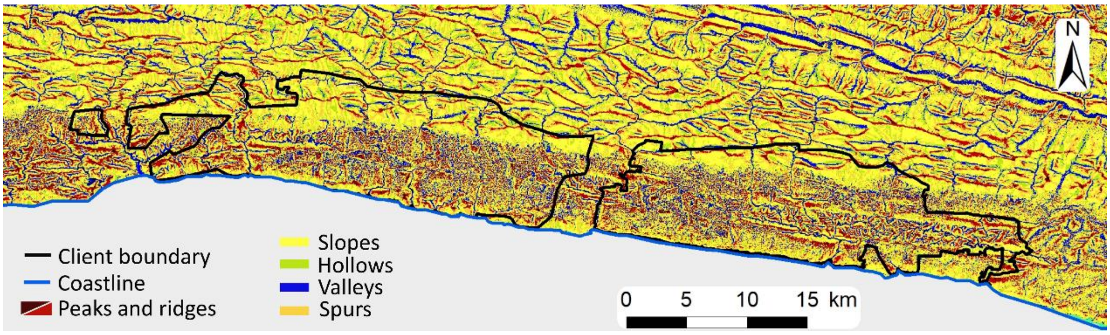

(b)

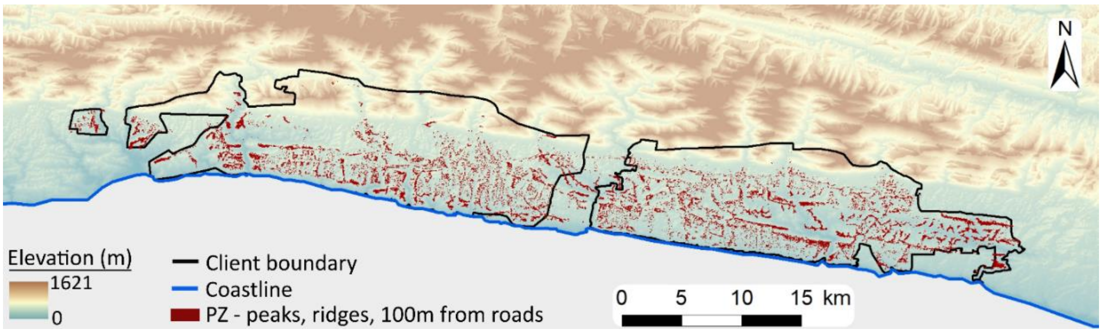

(c)

Figure 10 Process for identifying the PZ for the Southern Cape site-selection problem. (a) Client boundaries and terrain within 100 m of roads (in grey), considered for accessibility purposes. (b) Geomorphon classification of the Southern Cape terrain (only the most notably visible landforms are provided in the legend). (c) The final selection of candidate sites (final PZ).

Two CZs were considered, namely (a) 30 m above plantations with a 500 m buffer (immediate detection), and (b) 100 m above plantations with a 4 km buffer (secondary detection). Additionally, it was requested that the analysis complemented the cover achieved from a pre-selected tower, located slightly outside and to the east of the client's boundaries. This tower site was previously confirmed, meaning that certain portions of the smoke layers and buffers would already be covered by a camera there, so these covered areas were excluded from the remaining cover demand. In Figure 11 the CZ boundaries are shown, along with the location and cover achieved by the pre-determined camera, as well the coverage achieved by the ForestWatch-proposed site locations. The coverages were determined with a camera range of 8 km and proposed tower heights of 36 m, 24 m, 24 m, and 12 m, when moving from left to right in the images. The confirmed tower's coverage was determined at its proposed tower height of 24 m. It was determined that the four-tower layout proposed by ForestWatch experts would achieve cover of 58.9% and 45.8% with respect to the uncovered areas of CZ1 and CZ2, respectively. Noteworthy are the gaps in coverage that exist at 30 m in Figure 11(a) that are filled when the coverage of the towers are evaluated at 100 m in Figure 11(b). Clear examples exist to the southeastern corner of the cover achieved by the second-from-left tower, while also clearly visible to the south of the fourth-from-left tower. This shows that the CZs are not merely the client boundaries with extended buffer zones, but that the concept of determining smoke coverage at different heights above the terrain does indeed influence the coverage results.

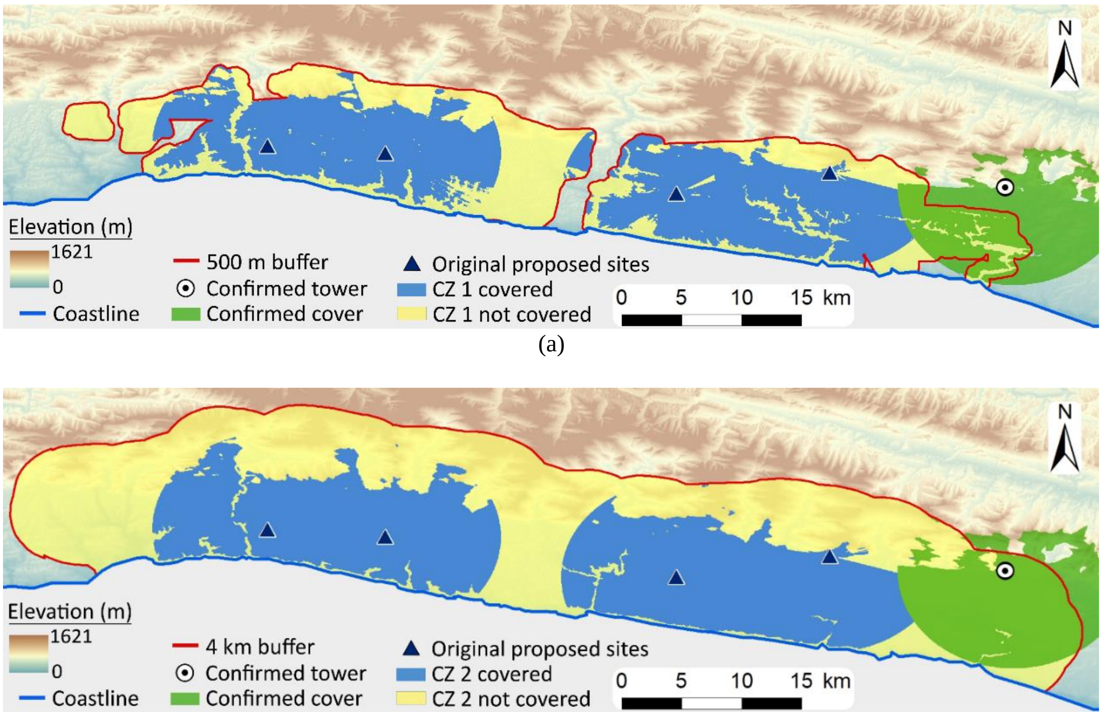

(a)

(b)

Figure 11 Site locations and cover achieved for the CWDS layout originally proposed by ForestWatch. Cover achieved with respect to (a) CZ1 and (b) CZ2.

3.3 System-site solution framework

As discussed before, both the use of heuristics and a weighted-sum approach have advantages and disadvantages. The use of heuristics is expected to continue forming the basis of the first optimisation stage, especially considering the general size of solution spaces considered in our problems. For the second stage, in which a weighted-sum approach is possible, we propose a combined approach – differing from our previous research in which the second stage involved either heuristics (Heyns et al., 2019) or ILP (Heyns et al., 2020), but not both. The benefits of a combined approach are numerous, and are demonstrated in the results presented later. Briefly, the weighted-sum approach may be used to determine, at the very least, the end-points of the Pareto front to provide an indication of its extent. We then employ both heuristics and the weighted-sum approach to determine solutions along the front, between the end-points. Instead of choosing one approach over the other, the weighted-sum approach can be used to approximate the general shape of the front, while the heuristic approach can be used to find numerous additional solutions between these points. Selected heuristic solutions may be proposed to decision-makers if their solution quality is considered acceptable when compared to the Pareto front's weighted-sum solutions, while weighted-sum solutions may, of course, also be proposed. An overview of the site-selection framework, divided into the GIS component and its two stages of the optimisation component, is provided in Figure 12.

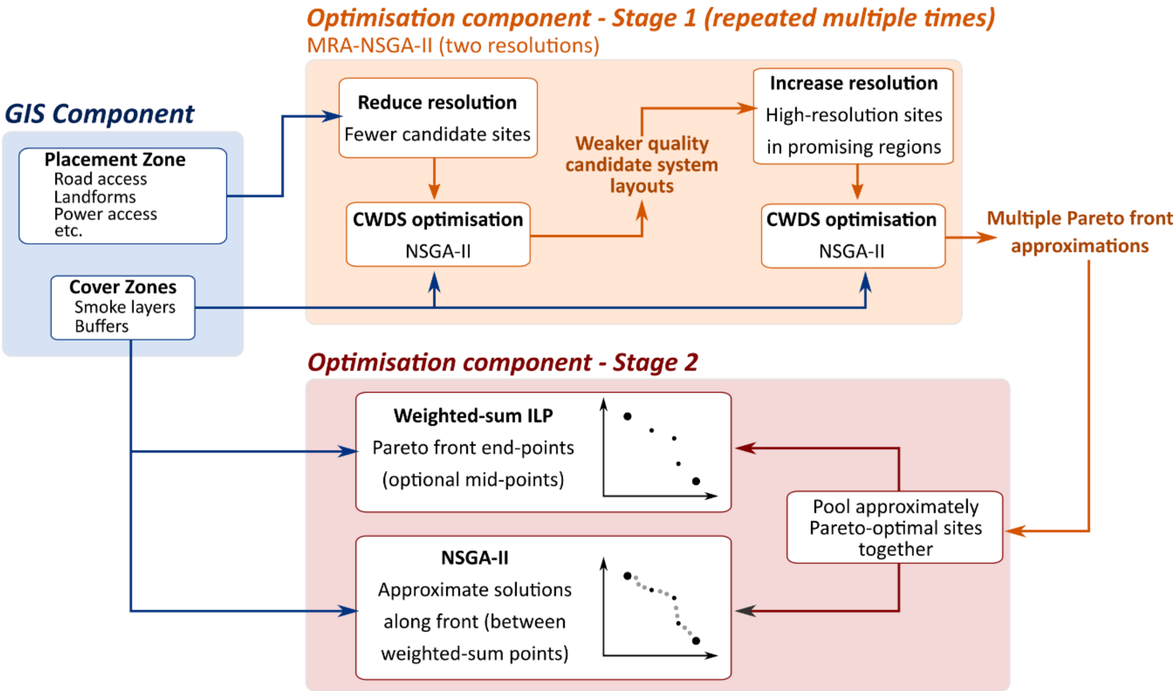

Figure 12 The CWDS tower site-selection optimisation framework, comprising a GIS component and two stages of optimisation components.

4 Results

4.1 Single-site alternative searches

Since three criteria were considered and two alternatives were sought for each, a total of six alternative sites were to be expected for each proposed site. However, many alternative sites were discovered with respect to more than one criterion – e.g. an alternative site offering the best total cover also being found to offer the best client cover. Generally, at least four alternatives were identified for each site. In the worst-case scenario, two sites were found as alternatives – one site achieved the best coverage with respect to each criterion, while the other achieved the second-best with respect to each.

The coverage results of the proposed sites and their alternatives are displayed in Table 1 and an example of the presentation of the results that were provided to decision makers is displayed in Figure 13. The alternative site locations were exported to be viewable in Google Earth, and each alternative's coverage values and viewsheds could be toggled on and off by clicking on its icon. Also viewable in the figure is the covered client area (shaded green) and the uncovered client area (shaded red), upon which the towers' coverage maps could be viewed and appraised.

The first six towers that are listed Table 1 (along the rows) are those for which the decision-makers chose one of our alternatives above their proposed sites. These selected alternatives are displayed in bold (sometimes identified as a superior alternative with respect to two criteria), and sites that were identified with respect to multiple criteria are indicated by asterisks (corresponding to the number of asterisks). Decision-makers explained that the six alternative sites were selected because of superior coverage and, in some instances, superior accessibility compared to the initially proposed sites. Regarding the remaining sites that the decision-makers chose to keep instead of choosing an alternative, in all instances these sites achieved inferior cover, but the decision-makers preferred them to alternatives due to either existing infrastructure or accessibility. The total hectares covered for each tower-criterion combination are displayed at the bottom of Table 1, and the total percentage improvement that was achieved by the alternatives compared to that achieved by the proposed towers with respect to each criterion are also displayed. The most significant value here is the almost 20% improvement in client cover that was determined to be possible with the alternatives.

Table 1 Summary of the coverage results of the single-site alternative search for 13 towers in the Mpumalanga province. Alternatives that were selected in favour of the proposed sites are displayed in bold, while sites that were identified with respect to multiple criteria are indicated by the number of asterisks.

	Proposed tower cover (hectares)			**Alternatives – total cover (hectares)**		**Alternatives – client cover (hectares)**		**Alternatives – new cover (hectares)**	
Tower name and number	***Total***	***Client***	***New***	***Best***	***Second***	***Best***	***Second***	***Best***	***Second***
Dundonald 1	9705	4401	2027	**12754***	11853**	4719	4715**	**2739***	2582
Zwalusnest 2	14089	2106	2095	14169*	**13784**	2111*	2060	2100*	2048
Blairmore 3	14365	10838	5983	16052	15825	13646	13159	**6204**	6144
Mkhondo 4	15699	1641	1546	17445	**17288***	2620	2425	**2054***	1969
Ridges 5	15502	5019	198	**15502**	14473	7197	6761	267	232
Derby 6	12501	7200	1801	**16073***	15973**	**9850***	9771**	1993	1980
Ntabanyama 7	10761	6248	847	12393*	11011**	7193*	6360**	891*	868**
Klipkopjie 8	14892	8826	608	16483*	16129	10194*	9738	721	678
Potgieterskeurs 9	15841	1170	1164	15949	15561	1211*	1195	1205*	1191
World's View 10	11028	10412	5148	12587*	12336**	11945*	11681**	5421	5117
Mac Mac 11	12575	6381	2208	13028*	11496**	6472	6311	2309*	1823**
Van Staden 12	14412	6169	941	14464	14112	7184	6747*	1206	1151*
Snymansbult 13	12742	6196	958	14091*	13802	7258	7121	1150	1133*
Total (ha)	174112	76607	25524	190990	183643	91600	88044	28260	26916
Improvement (%)				9.7	5.5	19.6	14.9	10.7	5.5

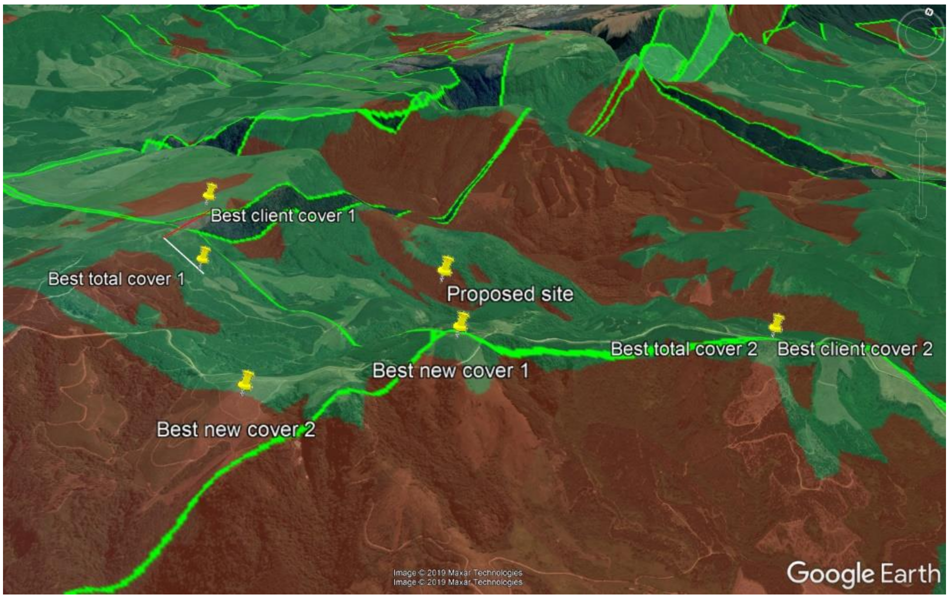

Figure 13 An illustration of solutions provided to decision makers in the single-site alternative search problems. Multiple alternatives were exported to view Google Earth and their coverage values and coverage maps could be toggled on and off by clicking on each alternative site (not illustrated in the figure). Coverage maps could be viewed and compared relative to the covered client areas (green shaded within green boundaries) and the uncovered client areas (the red shaded areas).

The files that were provided to ForestWatch and viewable as in *Figure 13* are available online (Heyns, 2020) and can be viewed with Google Earth Pro software. The user can toggle client areas, uncovered areas, and the thirteen sites and their alternatives' locations and viewsheds. Note that the Mkhondo alternative site search was expanded to 4 km instead of the 2 km used for all the other sites, due to specific technician request.

4.2 Southern Cape four-tower system

Twenty Pareto front approximations were performed with the MRA-NSGA-II, and their results and the final attainment front (the set of best solutions from all runs) are provided in Figure 14. A number of layouts from the attainment front were provided to decision makers and presented in a manner similar to that of the single-site selection process displayed in Figure 13 – layouts were exported to be viewable in Google Earth, along with proposed cover maps above client property. The solution that achieved the best cover with respect to CZ1 (i.e. the solution furthest to the right in Figure 13) was selected because of its coverage, but also because the sites were located in areas that were considered accessible and practical. The site locations and cover achieved with respect to the CZs by the two end-point solutions on the attainment front in Figure 13 – *i.e.* the two solutions that perform best with respect to each objective – are displayed in Figure 15(a) to (d) to illustrate the kind of trade-offs in site locations and cover that result from following a MO, multi-CZ solution approach. The results were obtained within four days – including data collection, processing, optimisation, analysis, and exports to visual presentation for decision-makers. CZ % achieved by these solutions?

As may be expected in a practical environment, ForestWatch experts decided to relocate the locations of some sites in the framework-determined solutions – the two eastern-most sites of the selected layout in this instance. The relocations are displayed in Figure 16 and were to improve accessibility for the site in (a) in the figure, while the site in (b) was moved to gain access to a stable power supply. These relocations resulted in minor changes to the coverage results and the relocated layout's objective function values are displayed in Figure 14 (the black cross). This only reduced the cover with respect to CZ1 by 4%, while the loss of cover with respect to CZ2 is negligible. Furthermore, compared to the layout proposed by decision makers, this final, relocated layout achieves an improvement of 9% in cover with respect to CZ1 and 10% with respect to CZ2, (not forgetting that the initial layout is evaluated at taller proposed tower heights, while the final layout is evaluated with 12-m towers).

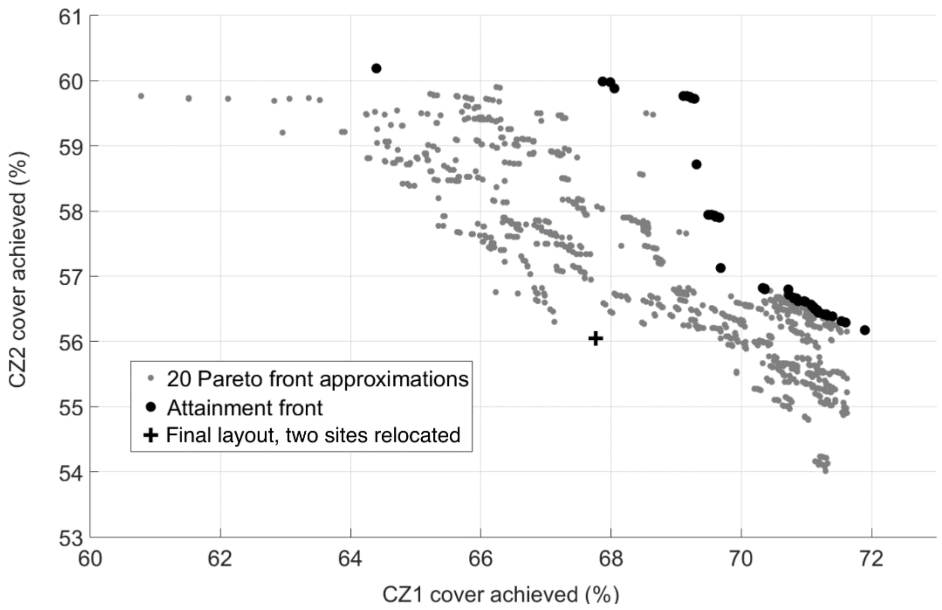

Figure 14 Pareto-front approximations of twenty heuristic runs (the grey dots) and the resulting attainment front (the black dots).

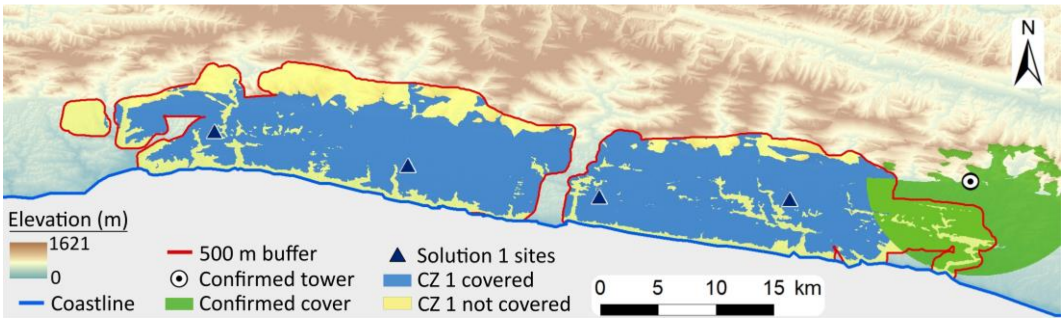

(a)

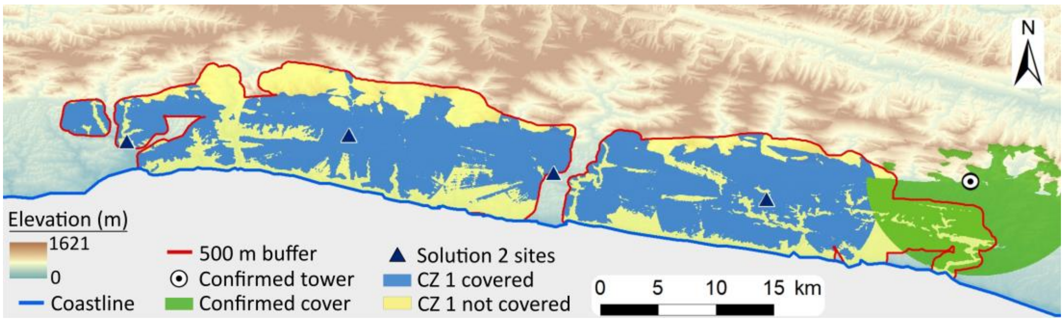

(b)

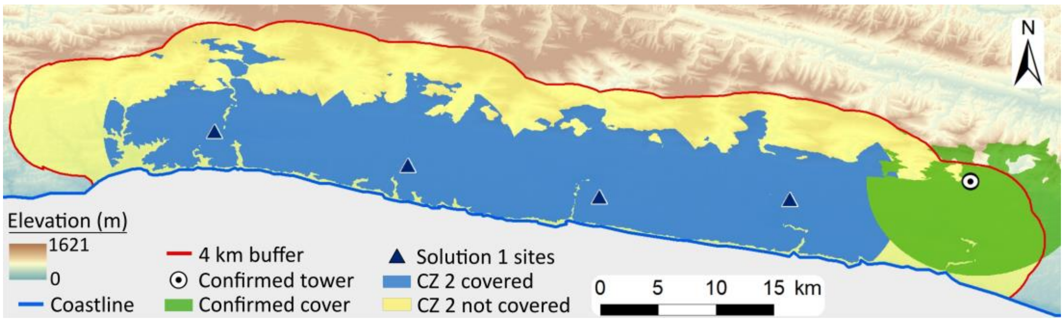

(c)

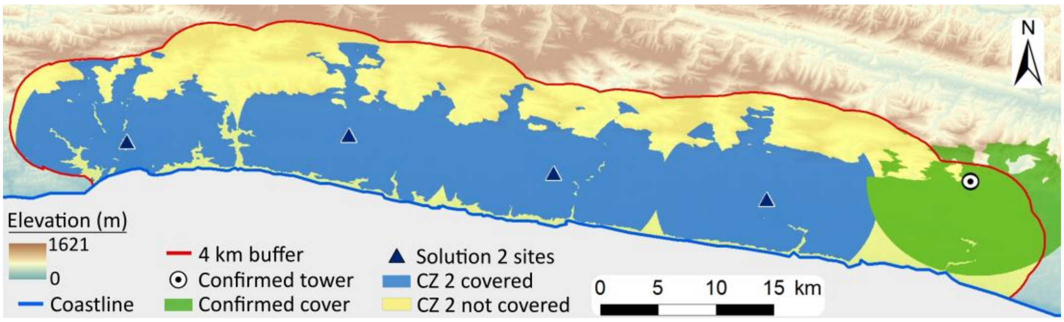

(d)

Figure 15 Site locations and cover achieved for the solutions from Figure 14 that achieve the best cover with respect to CZ1 (solution 1) and CZ2 (solution 2). Cover achieved with respect to CZ1 in (a) and (b), and CZ2 in (c) and (d).

(a) (b)

Figure 16 Relocations imposed by decision makers on the two eastern-most sites of the proposed layout in Figure 15(a) and (c) (solution 1). The relocation in (a) was for improved accessibility, while that in (b) was to gain access to a stable power supply.

Files for some of results presented here are available online (Heyns, 2020), and viewable in Google Earth Pro. These include the client areas, areas covered by the pre-specified tower, the CZs, and the site locations and viewsheds of the two solutions viewable in Figure 15. The two altered site locations as in Figure 16 are also provided.

4.3 Post-optimisation analysis

The Southern Cape CWDS site selection problem was performed within a limited timeframe and only the first stage of the optimisation component could be completed. The second stage of the optimisation component was performed afterwards to investigate what results could have been obtained if the full framework had been implemented, and to evaluate the quality of the solutions that had been proposed.

All the sites contained in the solutions in the twenty Pareto front approximations in Figure 14 were pooled together, resulting in a small PZ of 363 candidate sites. Ten additional runs of the NSGA-II were performed with this PZ as input and the attainment front achieved by these runs is provided in Figure 17 (the empty black markers). Furthermore, the same PZ was used as input to a weighted-sum ILP approach, with the following weights used for the two objectives presented in the format (CZ1, CZ2): (1.00,0.00), (0.75,0.25), (0.5,0.5), (0.25,0.75), (0.00, 1.00). The first and last weight sets effectively examine the optimal solution for a single CZ. The resulting solutions for this approach are shown in Figure 17 by the red markers. Of note is that there appears to be no real improvement in the quality of the solutions obtained by the first optimisation stage (the blue markers) and those from the second stage. The explanation for this is a smaller number of towers to place and a significantly smaller PZ than investigated in the Nelspruit problems (Heyns et al., 2019, 2020) in which the second optimisation stage demonstrated conspicuous improvement – here, 4 towers to place compared to 20, and 46 483 candidate sites compared to 741 813. The smaller computational complexity therefore results that the first, MRA-NSGA-II optimisation stage is able to obtain high quality site combinations.

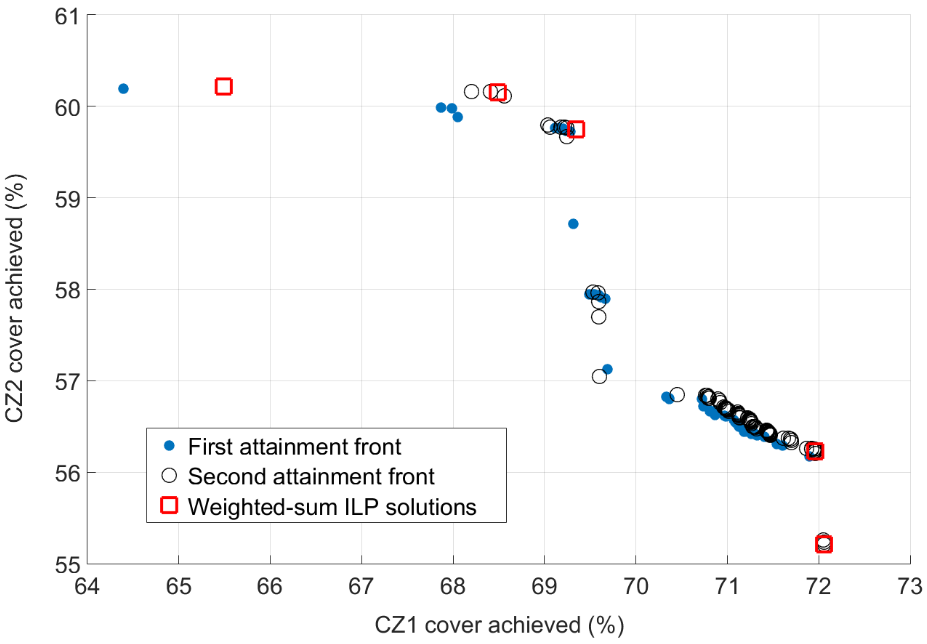

Figure 17 Pareto-front approximation attainment fronts obtained by additional heuristic and weighted-sum ILP runs from the second stage of the optimisation component.

The second-stage results demonstrate some of the disadvantages of both approaches. As previously stated and as may be seen here, evenly distributed ILP weights nevertheless result in an unevenly distributed weighted-sum Pareto front approximation. This is because of the changes in shape from concave (outer areas) to convex (the large gap area between two solutions), as has been described in the literature – see e.g. (Li and Yao, 2017). Nevertheless, the weighted-sum solutions provide us with an indication of where the true optimal front lies (at least for the 363 sites in the smaller PZ). This allows the quality of the heuristic-determined solutions to be better appraised. Concerning the solutions from the second-stage heuristic-determined front, it is clear that they exhibit similar solution quality compared to those offered by the weighted-sum front. The heuristic front includes a solution that matches one of the end-points of the weighted-sum front, but fails to reach the other – a known weakness of the heuristic approach (Kim et al., 2008). Furthermore, many solutions on the final heuristic front is observed (88 solutions). Such a large number of solutions is impractical for decision making purposes and many are clustered closely together and do not offer significant trade-offs in objective function values, nor tower site locations. Nevertheless, the front does discover numerous solutions otherwise overlooked in the gap of the weighted-sum front. Such solutions could offer coverage and tower site locations that may be of interest to decision makers in a practical environment, which can be overlooked if only a weighted-sum approach is followed.

The above observations illustrate how the approach of combining heuristic and weighted-sum analyses is capable of revealing important Pareto front characteristics and solution analysis to support CWDS decision-making and should certainly prevail in future problems.

5 Discussion

The results obtained for the single-site alternative searches and the system-optimisation problem were all well-received by decision-makers. The alternative search framework provided decision-makers with practical solutions which they could consider, and they appreciated the manner in which they could compare sites in Google Earth and visually display different site viewsheds. Unfortunately, ForestWatch was not awarded the most recent contract for wildfire detection in the region to build on their existing towers in operation. Despite the fact that the towers identified here may not be constructed, the purpose of our framework is to determine sites that can help ForestWatch decision-makers in selecting final sites, in real-world problems. This was achieved because the sites that were determined with the aid of our framework were those that were confirmed to be built if the contract was awarded.

The system-site selection framework was applied to solve a bi-objective CWDS placement problem in the Southern Cape – with the maximisation of coverage of each CZ as an objective – using MO optimisation solution approaches. Feedback from decision makers involved in this problem was that the solutions allowed them to spend their time on refining site locations, as opposed to performing rigorous practical site searches with no starting point at all. Furthermore, the coverage maps and the coverage values in hectares that are exported to be easily viewed in Google Earth make it possible to easily analyse and compare – and the nature of the proposed data and their presentation give the clients assurance that best efforts have been made in finding optimal sites. At the time of writing, completion of the contract in the area is still pending and has been delayed by the client due to economic factors, further exacerbated by Covid-19. Nevertheless, the goal of providing practical tower sites by the framework to aid decision-makers has been achieved. The sites proposed by our framework as presented in this paper – and agreed upon between ForestWatch and the client – will be used as soon as (and if) the contract finally proceeds.

In future problems it is expected that additional (or alternative) objectives may be considered. One example is that of proximity to power supplies. In the Southern Cape exercise, proximity from power was not considered as a constraint in determining PZ sites because of typical problems that ForestWatch have experienced with power outages and their preference for installing towers with solar power supplies. As was observed during solution analysis and selection, however, experts decided to move one tower from their preferred layout to another because of access to power. It is therefore expected that power proximity should not be used as a limiting constraint on PZ identification – instead, the minimisation of distance to power supplies should be considered as an additional *objective*. Decision-makers may then consider distance to power along with covering objectives in their analysis of Pareto-front solutions. An alternative approach would be to continue determining solutions without power supply considerations and, instead, perform local "fine-tuning" of solutions. This can be performed by automated search algorithms which determining the closest power supply point to each tower in a solution and if it is within a suitable distance, the site in the solution may be exchanged for the nearest suitable site to the power supply, and the effect on the coverage results determined. If the changes are within an acceptable threshold, then the new site may be kept.

Additional CZs which may be considered in future problems include certain priority areas within the larger client area, e.g. areas around key infrastructure points such as power plants and chemical storage facilities, or areas that are historically fire-prone. In such instances, a priority CZ is simply added as an additional covering objective and the problem solved as usual by the MO optimisation framework. Furthermore, alternative or additional objectives that may be considered include those that were investigated in the single-site searches, namely 1) total coverage, 2) client coverage, and 3) new client coverage – as opposed to smoke layers at different heights.

A set of solutions that is diverse in respect of objective function values and tower site locations are desired for decision makers. It is, however, possible that attainment fronts consisting of an undesirably large number of solutions may be returned, as was observed during the Southern Cape post-optimisation analysis. Many of the solutions appear in "clusters" and offer negligible trade-offs in terms objective function values and facility site locations. In these instances, reduction techniques may be considered to filter the Pareto front to an acceptable number of solutions. Such techniques include those that are performed in objective function space, such as the epsilon-grid method (Mavrotas, 2009), or those performed in physical solution space, such as site-proximity de-clustering (Heyns, 2016) – a combination of such techniques also merits investigation.

6 Conclusion

The development of two comprehensive CWDS tower-site selection optimisation frameworks for single-site alternative searches and system-site optimisation for implementation in vast, unknown territories has been described and practically applied in South Africa. The main aims of the framework are to determine multiple candidate CWDS layouts within short timeframes with minimal user input. First, the single-site selection problem provided a foundation for the development of the GIS component, and also contributed to the process of visualising solutions to decision makers. Numerous alternative sites were found for thirteen sites proposed by ForestWatch, and six of their initially proposed sites were discarded for one of our proposals. Second, in the Southern Cape region, the framework obtained numerous superior solutions as alternatives to a four-tower layout proposed by ForestWatch (which required weeks of speculation and planning to determine). Our alternatives were determined by the framework within four days. Multiple proposed system layouts and coverage maps were presented to decision-makers, who selected one of our proposed solutions due to its superior cover and practical tower site locations. The layouts obtained by the optimisation framework were found to significantly outperform the initial layout with respect to both covering objectives – despite the optimisation solutions being limited to 12-m tower heights while the proposed system had an average tower height of 24 m. The fact that the installation cost of a 12-m tower is less than half that of a 24-m tower is an indication of the potential cost savings that may be achieved by the optimisation approach.

Going forward, the frameworks are planned for implementation in future ForestWatch site-selection problems, with numerous opportunities for improvement as described in the discussion.

7 Acknowledgements

The authors would like to express their sincere thanks to Mr Dennis Lawrie, Mr Dave Kuhl, Mr Adrian Daniel and Mr Gareth Perks of ForestWatch for their valuable discussions, suggestions and timely data provision during the development of our framework.

8 References

Alp, O., Erkut, E., Drezner, Z., 2003. An Efficient Genetic Algorithm for the p-Median Problem. Annals of Operations Research 122, 21–42. https://doi.org/10.1023/A:1026130003508

Bao, S., Xiao, N., Lai, Z., Zhang, H., Kim, C., 2015. Optimizing watchtower locations for forest fire monitoring using location models. Fire Safety Journal 71, 100–109. https://doi.org/10.1016/j.firesaf.2014.11.016

Church, R., ReVelle, C., 1974. The maximal covering location problem. Papers of the Regional Science Association 32, 101–118.

Cohon, J.L., 1978. Multiobjective Programming and Planning. Academic Press, New York (NY).

Deb, K., Pratap, A., Agarwal, S., Meyarivan, T., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 6, 182–197. https://doi.org/10.1109/4235.996017

Di Stefano, M., Mayer, L., 2018. An Automatic Procedure for the Quantitative Characterization of Submarine Bedforms. Geosciences 8, 28. https://doi.org/10.3390/geosciences8010028

Djurdjevac Conrad, N., Helfmann, L., Zonker, J., Winkelmann, S., Schütte, C., 2018. Human mobility and innovation spreading in ancient times: a stochastic agent-based simulation approach. EPJ Data Science 7. https://doi.org/10.1140/epjds/s13688-018-0153-9

Eugenio, F.C., Rosa dos Santos, A., Fiedler, N.C., Ribeiro, G.A., da Silva, A.G., Juvanhol, R.S., Schettino, V.R., Marcatti, G.E., Domingues, G.F., Alves dos Santos, G.M.A.D., Pezzopane, J.E.M., Pedra, B.D., Banhos, A., Martins, L.D., 2016. GIS applied to location of fires detection towers in domain area of tropical forest. Science of The Total Environment 562, 542–549. https://doi.org/10.1016/j.scitotenv.2016.03.231

Fonseca, C.M., Fleming, P.J., 1993. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization, in: Proceedings of the Fifth International Conference on Genetic Algorithms. pp. 416–423.

Forsythe, G., Le Maitre, D., van den Dool, R., Walls, R., Pharoah, R., Fortune, G., 2019. The Knysna Fires of 2017: Learning from this disaster [WWW Document]. URL https://www.polity.org.za/article/the-knysna-fires-of-2017-learning-from-this-disaster-2019-06-07 (accessed 12.18.19).

Franklin, W.R., 2002. Siting observers on terrain, in: Richardson, D.E., van Oosterom, P. (Eds.), Advances in Spatial Data Handling. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 109–120. https://doi.org/10.1007/978-3-642-56094-1_9

Franklin, W.R., Clark, R., 1994. Higher isn't necessarily better: Visibility algorithms and experiments, in: Advances in GIS Research: Sixth International Symposium on Spatial Data Handling. Taylor & Francis, pp. 751–770.

Harmon, B.A., Petrasova, A., Petras, V., Mitasova, H., Meentemeyer, R., 2018. Tangible topographic modeling for landscape architects. International Journal of Architectural Computing 16, 4–21. https://doi.org/10.1177/1478077117749959

Heyns, A.M., 2020. Google Earth client and viewshed files for ITOR paper [WWW Document]. URL http://www.andriesheyns.com/ITOR.zip (accessed 11.23.20).

Heyns, A.M., 2016. A multi-objective approach towards geospatial facility location (Author PhD Thesis). Stellenbosch University, Stellenbosch.

Heyns, A.M., du Plessis, W., Curtin, K.M., Kosch, M., Hough, G., 2020. Analysis and exploitation of landforms for improved optimisation of camera-based wildfire detection systems. Fire Technology Under Review.

Heyns, A.M., du Plessis, W., Kosch, M., Hough, G., 2019. Optimisation of tower site locations for camera-based wildfire detection systems. International Journal of Wildland Fire 28, 651–665.

Heyns, A.M., Van Vuuren, J.H., 2018. Multi-type, multi-zone facility location. Geographical Analysis 32, 3–31. https://doi.org/10.1111/gean.12131

Heyns, A.M., van Vuuren, J.H., 2016. A multi-resolution approach towards point-based multi-objective geospatial facility location. Computers, Environment and Urban Systems 57, 80–92. https://doi.org/10.1016/j.compenvurbsys.2016.01.007
Heyns, A.M., Van Vuuren, J.H., 2015. An evaluation of the effectiveness of observation camera placement within the MeerKAT radio telescope project. The South African Journal of Industrial Engineering 26, 1–10. https://doi.org/10.7166/26-2-1216
Hough, G., 2007. Vision systems for wide area surveillance: ForestWatch -- a long-range outdoor wildfire detection system, in: TASSIE FIRE Conference Proceedings. Presented at the The Tassie Fire Conference, Hobart, Tasmania.
Jasiewicz, J., Stepinski, T.F., 2013. Geomorphons - a pattern recognition approach to classification and mapping of landforms. Geomorphology 182, 147–156. https://doi.org/10.1016/j.geomorph.2012.11.005
Jia, H., Ordóñez, F., Dessouky, M.M., 2007. Solution approaches for facility location of medical supplies for large-scale emergencies. Computers & Industrial Engineering 52, 257–276.
Khan, S.A., Rehman, S., 2013. Iterative non-deterministic algorithms in on-shore wind farm design: A brief survey. Renewable and Sustainable Energy Reviews 19, 370–384. https://doi.org/10.1016/j.rser.2012.11.040
Kim, K., Murray, A.T., Xiao, N., 2008. A multiobjective evolutionary algorithm for surveillance sensor placement. Environment and Planning B: Planning and Design 35, 935–948. https://doi.org/10.1068/b33139
Kim, Y.H., Rana, S., Wise, S., 2004. Exploring multiple viewshed analysis using terrain features and optimisation techniques. Computers & Geosciences 30, 1019–1032. https://doi.org/10.1016/j.cageo.2004.07.008
Kwong, W.Y., Zhang, P.Y., Romero, D., Moran, J., Morgenroth, M., Amon, C., 2014. Multi-objective wind farm layout optimization considering energy generation and noise propagation with NSGA-II. Journal of Mechanical Design 136, 1–10.
Lee, J., 1994. Digital Analysis of Viewshed Inclusion and Topgraphic Features on Digital Elevation Models. Photogrammetric Engineering & Remote Sensing 60, 451–456.
Li, M., Yao, X., 2017. What Weights Work for You? Adapting Weights for Any Pareto Front Shape in Decomposition-based Evolutionary Multi-Objective Optimisation. arXiv:1709.02679 [cs].
Luo, W., Liu, C.C., 2018. Innovative landslide susceptibility mapping supported by geomorphon and geographical detector methods. Landslides 15, 465–474. https://doi.org/10.1007/s10346-017-0893-9
Machairas, V., Tsangrassoulis, A., Axarli, K., 2014. Algorithms for optimization of building design: A review. Renewable and Sustainable Energy Reviews 31, 101–112. https://doi.org/10.1016/j.rser.2013.11.036
Marler, R.T., Arora, J.S., 2010. The weighted sum method for multi-objective optimization: new insights. Structural and Multidisciplinary Optimization 41, 853–862. https://doi.org/10.1007/s00158-009-0460-7
Marler, R.T., Arora, J.S., 2004. Survey of multi-objective optimization methods for engineering. Structural and Multidisciplinary Optimization 26, 369–395. https://doi.org/10.1007/s00158-003-0368-6
Martell, D.L., 2015. A Review of Recent Forest and Wildland Fire Management Decision Support Systems Research. Current Forestry Reports 1, 128–137. https://doi.org/10.1007/s40725-015-0011-y
Mavrotas, G., 2009. Effective implementation of the epsilon-constraint method in Multi-Objective Mathematical Programming problems. Applied Mathematics and Computation 213, 455–465.
Meunier, H., Talbi, E., Reininger, P., 2000. A multiobjective genetic algorithm for radio network optimization, in: Proceedings of the 2000 Congress on Evolutionary Computation. IEEE, pp. 317–324.

Misthos, L.M., Nakos, B., Krassanakis, V., Menegaki, M., 2018. The effect of topography and elevation on viewsheds in mountain landscapes using geovisualization. International Journal of Cartography 5, 44–66. https://doi.org/10.1080/23729333.2018.1477569

Murray, A.T., Kim, K., Davis, J.W., Machiraju, R., Parent, R., 2007. Coverage optimization to support security monitoring. Computers, Environment and Urban Systems 31, 133–147. https://doi.org/10.1016/j.compenvurbsys.2006.06.002

Nagy, G., 1994. Terrain visibility. Computers & Graphics 18, 763–773. https://doi.org/10.1016/0097-8493(94)90002-7

O'Sullivan, D., Turner, A., 2001. Visibility graphs and landscape visibility analysis. International Journal of Geographical Information Science 15, 221–237. https://doi.org/10.1080/13658810151072859

Owen, S.H., Daskin, M.S., 1998. Strategic facility location: A review. European journal of operational research 111, 423–447.

Purshouse, R.C., Fleming, P.J., 2003. Evolutionary many-objective optimisation: an exploratory analysis, in: The 2003 Congress on Evolutionary Computation, 2003. Presented at the The 2003 Congress on Evolutionary Computation, 2003. CEC '03, pp. 2066–2073. https://doi.org/10.1109/CEC.2003.1299927

Raisanen, L., Whitaker, R.M., 2005. Comparison and evaluation of multiple objective genetic algorithms for the antenna placement problem. Mobile Networks and Applications 10, 79–88.

Rana, S., 2003. Fast approximation of visibility dominance using topographic features as targets and the associated uncertainty. Photogrammetric Engineering & Remote Sensing 69, 881–888. https://doi.org/10.14358/PERS.69.8.881

Rego, F.C., Catry, F.X., 2006. Modelling the effects of distance on the probability of fire detection from lookouts. Int. J. Wildland Fire 15, 197–202. https://doi.org/10.1071/WF04016

ReVelle, C.S., Eiselt, H.A., 2005. Location analysis: A synthesis and survey. European Journal of Operational Research 165, 1–19. https://doi.org/10.1016/j.ejor.2003.11.032

Romero, B.E., Clarke, K.C., 2018. Exploring uncertainties in terrain feature extraction across multi-scale, multi-feature, and multi-method approaches for variable terrain. Cartography and Geographic Information Science 45, 381–399. https://doi.org/10.1080/15230406.2017.1335235

Schroeder, D., 2005. Operational trial of the ForestWatch wildfire smoke detection system. Advantage 6, 1–7.

South African Government, 2009. Forestry | South African Government [WWW Document]. URL https://www.gov.za/about-sa/forestry (accessed 12.14.19).

Strydom, S., Savage, M.J., 2016. A spatio-temporal analysis of fires in South Africa. South African Journal of Science 112, 1–8. https://doi.org/10.17159/sajs.2016/20150489

Tabik, S., Zapata, E.L., Romero, L.F., 2013. Simultaneous computation of total viewshed on large high resolution grids. International Journal of Geographical Information Science 27, 804–814. https://doi.org/10.1080/13658816.2012.677538

Tanergüçlü, T., Maraş, H., Gencer, C., Aygüneş, H., 2012. A decision support system for locating weapon and radar positions in stationary point air defence. Information Systems Frontiers 14, 423–444. https://doi.org/10.1007/s10796-010-9269-6

Tong, D., Murray, A., Xiao, N., 2009. Heuristics in spatial analysis: A genetic algorithm for coverage maximization. Annals of the Association of American Geographers 99, 698–711.

Xia, J., Curtin, K.M., Huang, J., Wu, D., Xiu, W., Huang, Z., 2019. A carpool matching model with both social and route networks. Computers, Environment and Urban Systems 75, 90–102. https://doi.org/10.1016/j.compenvurbsys.2019.01.008

Xiao, N., Bennett, D.A., Armstrong, M.P., 2002. Using evolutionary algorithms to generate alternatives for multiobjective site-search problems. Environment and Planning A 34, 639–656. https://doi.org/10.1068/a34109

Yamani Douzi Sorkhabi, S., Romero, D.A., Yan, G.K., Gu, M.D., Moran, J., Morgenroth, M., Amon, C.H., 2016. The impact of land use constraints in multi-objective energy-noise wind farm layout optimization. Renewable Energy 85, 359–370. https://doi.org/10.1016/j.renene.2015.06.026

Yao, J., Zhang, X., Murray, A.T., 2018. Spatial Optimization for Land-use Allocation: Accounting for Sustainability Concerns. International Regional Science Review 41, 579–600. https://doi.org/10.1177/0160017617728551

Zhang, F., Zhao, P., Thiyagalingam, J., Kirubarajan, T., 2019. Terrain-influenced incremental watchtower expansion for wildfire detection. Science of The Total Environment 654, 164–176. https://doi.org/10.1016/j.scitotenv.2018.11.038

Zitzler, E., Laumanns, M., Bleuler, S., 2004. A Tutorial on Evolutionary Multiobjective Optimization, in: Gandibleux, X., Sevaux, M., Sörensen, K., T'kindt, V. (Eds.), Metaheuristics for Multiobjective Optimisation. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 3–37. https://doi.org/10.1007/978-3-642-17144-4_1