

Metaheuristic Approaches to the Placement of Suicide Bomber Detectors

Carlos Cotta · José E. Gallardo

Received: date / Accepted: date

Abstract Suicide bombing is an infamous form of terrorism that is becoming increasingly prevalent in the current era of global terror warfare. We consider the case of targeted attacks of this kind, and the use of detectors distributed over the area under threat as a protective countermeasure. Such detectors are non-fully reliable, and must be strategically placed in order to maximize the chances of detecting the attack, hence minimizing the expected number of casualties. To this end, different metaheuristic approaches based on local search and on population-based search (such as a Hill Climber, different Greedy Randomized Adaptive Search Procedures, an Evolutionary Algorithm and several Estimation of Distribution Algorithms) are considered and benchmarked against a powerful greedy heuristic from the literature. We conduct an extensive empirical evaluation on synthetic instances featuring very diverse properties. Most metaheuristics outperform the greedy algorithm, and a hill-climber is shown to be superior to remaining approaches. This hill-climber is subsequently subject to a sensitivity analysis to determine which problem features make it stand above the greedy approach, and is finally deployed on a number of problem instances built after realistic scenarios, corroborating the good performance of the heuristic.

Keywords Counter-terrorism · Suicide Bombing · Optimal Detector Placement · Greedy Heuristics · Metaheuristics

Authors acknowledge support from Spanish Ministry of Economy and Competitiveness and European Regional Development Fund (FEDER) under project EphemeCH (TIN2014-56494-C4-1-P).

Carlos Cotta · José E. Gallardo
Department *Lenguajes y Ciencias de la Computación*, Universidad de Málaga,
Campus de Teatinos, 29071 - Málaga, Spain.
Tel.: +34 95 213 7158
Fax: +34 95 213 1397
E-mail: {ccottap,pepeg}@lcc.uma.es

1 Introduction

At the time of writing this paper, the Atatürk Airport in Istanbul was subject to a terrorist attack whose perpetrators carried automatic weapons and explosive belts. As a result of the shooting and the detonation of the suicide bombs, 44 people were killed (in addition to the three terrorists) and more than 200 were injured (Şenses Kurukız, 2016). A few days afterwards, a suicide car bombing killed 292 people and left more than 200 people wounded in a shopping district of Baghdad, in what constituted the deadliest single bomb attack in Iraq since 2007 (AFP, 2016). These are only two recent examples of suicide bombings, a form of terrorism that not only causes about four times more casualties than other kinds of terrorist attack (Rand Corporation, 2009), cf. (Hoffman, 2003) but also instills a sense of fear in the society as a whole that undermines public confidence in the authorities and contributes to subjugate those living under this threat (Hoffman, 2003). As a result of the relative inexpensiveness –they do not require escaping logistics nor sophisticated equipment for remote operation– and effectiveness –a combination of high lethality and target accuracy, with a statistically significant higher kill-to-wounded ratio than non-suicide attacks, see (Edwards et al, 2016)– of this kind of attacks, they have become increasingly prevalent; see Figure 1. Indeed, they are not just very frequent in conflicting areas such as the Middle East –1,192 attacks in the period 1982-2015, only counting those conducted using explosive belts (Chicago Project on Security and Terrorism (CPOST), 2016)– but also constitute a global threat that has caused nearly 50,000 deaths worldwide in the 21st Century, with an average of about 10 fatalities and 24 non-fatal casualties per attack (counting all suicide attacks – data from Chicago Project on Security and Terrorism (CPOST) (2016)).

In response to this ongoing threat, security forces and intelligence agencies are strengthening their efforts. Obviously, the nature of this kind of attack makes it infeasible to rely on standard deterrence measures based on direct retaliation. Hence, other members of the terrorist network providing intellectual or logistic support to the attacker (and in general any assets valued by the latter) are targeted instead (Kroenig and Pavel, 2012). In any case, this does not fully deter this kind of attacks, e.g., see (Carson, 2017), so they must be also fought against by trying to deny their benefits, either at a strategic level (ensuring that the ultimate goals of the terrorists will not be achieved even if the attack was successful) or at a tactical level (trying to make the attacks unsuccessful). Heightened security measures in airports, government buildings or military installations are examples of this latter kind of counterterrorism measures.

A potential problem of these explicit, partly invasive security measures is that they cannot be readily exported to other kind of environments also susceptible to terror attacks (e.g., placing airport-like scanners surrounding a city square or busy shopping street, or frisking every passer-by is out of question). More subtle surveillance can be accomplished by using chemical or biological sensors tailored to detect traces of explosives in their proximity (NRC, 2004; Singh, 2007; Yinon, 2007; Caygill et al, 2012; Gares et al, 2016). In this sense, Kaplan and Kress (2005) studied the deployment of such sensors in urban areas subject to pedestrian suicide-bombing attacks. They analyzed different scenarios (an urban grid of streets and a large plaza or park) and concluded that, while such sensors are not likely effective against random




Fig. 1: The bar chart shows the number of casualties (dead and injured) in suicide attacks in the period 2001-2015 (total year figures on top of each bar, scale on the left). The black curve indicates the number of suicide attacks in this time frame (scale on the right). Source: own elaboration based on data from (Chicago Project on Security and Terrorism (CPOST), 2016).

attacks, they can play an important role in the defense of known targets. Building on this work, Nie et al (2007) considered the case of a threat area with known targets on which a number of non-fully reliable sensors are deployed. They proposed a branch-and-bound algorithm (BnB) and a greedy constructive heuristic to determine the location of these sensors so as to minimize the casualties. The BnB algorithm does not scale well with problem size, but the greedy heuristic is relatively effective (although it will not guarantee finding the optimal solution in general). This same approach has been recently applied by Yan and Nie (2016) in order to tackle a hypothetical threat scenario where terrorists would utilize a small vessel to attack a maritime target. A real-world case study corresponding to New York Harbor was computationally addressed and results obtained for this particular instance by BnB and greedy constructive algorithms were compared and analyzed.

Following these previous results, the problem is here approached for the first time to the best of our knowledge from the point of view of iterative heuristics and metaheuristic techniques. We propose several algorithms based on local search, greedy randomized adaptive search procedures and evolutionary algorithms, and analyze their performance, comparing it to the greedy approach previously mentioned on an extensive collection of synthetic and real scenarios. Before presenting the algorithms considered and the experimental setup, let us firstly define formally the problem tackled. This is done in next section.

2 Problem Description

Let us assume the goal is to protect a certain threat area or scenario which we can model as a rectangular grid $A = \{A_{ij}\}_{m \times n}$. Each cell in this grid represents a small square subarea which can be *blocked* if there is some physical obstruction (a wall, a monument, etc.) that precludes its being traversed by the attacker, or *unblocked* if it can be freely accessed. Some of these latter unblocked cells can be also *objectives*, that is, cells that contain threatened individuals who may be targeted by the attacker. To carry on this threat, the attacker can enter the scenario through some *entrances*, namely specific unblocked cells typically placed in the boundaries of the grid. From any of those entrances, the attacker will try to reach any of the objectives following always a shortest path. Following (Nie et al, 2007), the attacker is able to move in a straight line from the center of a grid cell to the center of any other one for which this line does not intersect with a blocked cell. Thus, the path followed by the attacker will be composed by a number of straight segments¹, and can be computed by running classical algorithms –such as, e.g., Dijkstra’s algorithm or A*– on a weighted complete graph $G(V, E, w)$, where

- $V = \{\langle i, j \rangle \mid A_{ij} \text{ is unblocked}\}$, i.e., a vertex per unblocked cell,
- $E = V \times V$, and
- $w : E \rightarrow \mathbb{R}$ is defined as

$$w(\langle \langle i_1, j_1 \rangle, \langle i_2, j_2 \rangle \rangle) = \zeta \sqrt{(i_1 - i_2)^2 + (j_1 - j_2)^2}, \quad (1)$$

(where ζ is the side of each grid cell) if there is no obstacle in the straight path among both cells, and $w(E) = \infty$ otherwise.

Figure 2 shows an example 8×8 scenario with 8 entrances, 6 blocked cells and 2 objectives. Notice how paths try to follow always a straight line towards the objective, taking small detours if there are obstacles on the way.

With this setting in mind, the problem under consideration amounts to placing some *detectors* in different unblocked cells of the grid. These detectors are perfectly concealed from the attacker, but they are not fully reliable and will only detect attackers that travel within a certain detection radius τ of the detector with a probability p that depends on the length of the attacker’s path inside this detection area (a circle of radius τ centered in the detector’s location). Detection does not imply neutralization of the attacker though. It only amounts to firing some alarm that elicits response from some enforcing agents, and that can lead to effective negation of the attack only if the attacker is not yet very close to the objective, and with some probability θ . Let us formalize the whole process in the following.

¹ Note that the attacker following a non-shortest path would in principle imply a larger exposure to detection and hence this assumption provides a baseline for evaluation, and is used too in related contexts such as routing military ground operations (Kim et al, 2016). Also, according to the literature on modeling realistic human walking paths, there seems to be a governing principle which is minimizing the time-derivative of the curvature of the path (Arechavaleta et al, 2008). This mostly affects those cases in which a turn has to be made, resulting in clothoid-like trajectories. Notice however that this effect may be minor in light of factors such as the discretization of the map and the presence of obstacles (such as, e.g., street walls) constraining the trajectory.




Fig. 2: An 8×8 scenario. Green cells correspond to entrances, blue ones to objectives and gray cells represent blocked locations. Lines are shortest paths from each entrance to each objective (note that paths would be extended backwards should a detector be placed on any of the entrances). The dimension of each cell is $10\text{m} \times 10\text{m}$, and the attacker must be neutralized at least 10m away from an objective (we assume that the terrorist moves at a speed of 1m/s and that aborting the attack requires no more than 10 seconds).

Let ε and ϕ be the number of entrances and objectives respectively. For each entrance e_i , $1 \leq i \leq \varepsilon$, and objective o_j , $1 \leq j \leq \phi$ there is a path P_{ij} (we assume for simplicity that the shortest path is unique – considering multiple paths does not fundamentally affect the analysis below). The attacker will pick an entrance e_i and an objective o_j with probability γ_{ij} (obviously, $\sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} = 1$) and will subsequently move along P_{ij} at speed v . Neutralizing the attack once detected takes some time t_n – at least 10s according to Kaplan and Kress (2005) – and hence once the attacker is at distance vt_n from the objective (distance measured along the path P_{ij}) no effective neutralization is possible. Therefore, let us define \bar{P}_{ij} as the portion of the path P_{ij} outside this “dead” zone, in which timely detection is still possible.

Now, let δ be the number of detectors in the scenario. A certain detector d_k , $1 \leq k \leq \delta$, will timely detect an attacker traveling through P_{ij} with probability

$$p_{ijk} = 1 - \exp(-\eta l_{ijk}) \quad (2)$$

where l_{ijk} is the length of the segment of \bar{P}_{ij} within the circle of radius τ centered at d_k , and $\eta > 0$ is a parameter (the detector’s instantaneous detection rate). It follows that the larger this segment, the larger the detection probability (and obviously $p_{ijk} = 0$ if $l_{ijk} = 0$, i.e., if the path is outside the detection radius of d_k) – see Figure 3 for an example. Following (Nie et al, 2007), detectors are assumed to work independently of one another, and therefore the total probability of non-detection \tilde{D}_{ij} if the attacker




Fig. 3: A shortest path going from entrance e_i to objective o_j . Areas monitored by detectors d_k and d'_k are those enclosed by circumferences. The segments of this path detected by both detectors are shown with a dotted line and denoted by l_{ijk} and $l_{ijk'}$.

follows P_{ij} is

$$\tilde{D}_{ij} = \prod_{k=1}^{\delta} (1 - p_{ijk}) = \prod_{k=1}^{\delta} \exp(-\eta l_{ijk}) = \exp(-\eta \sum_{k=1}^{\delta} l_{ijk}) \quad (3)$$

In case of non-detection, the attacker will reach the objective o_j causing a number of casualties C_j . This will be also the case, should the attacker be timely detected but not effectively neutralized (an event that will happen with probability $1 - \theta$). The expected number of casualties W_{ij} for this particular path will thus be

$$W_{ij} = \tilde{D}_{ij} C_j + (1 - \tilde{D}_{ij})(1 - \theta) C_j = C_j [\tilde{D}_{ij} \theta + (1 - \theta)] \quad (4)$$

The total number of casualties will take into account all possible paths P_{ij} the attacker can take, that is, between any entrance e_i and objective o_j (and recall each of these is picked with probability γ_j):

$$W = \sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} W_{ij} = \sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} C_j [\tilde{D}_{ij} \theta + (1 - \theta)] = \quad (5)$$

$$= \sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} C_j (1 - \theta) + \sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} C_j \tilde{D}_{ij} \theta = \quad (6)$$

$$= (1 - \theta) \sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} C_j + \theta \sum_{i=1}^{\varepsilon} \sum_{j=1}^{\phi} \gamma_{ij} C_j \exp(-\eta \sum_{k=1}^{\delta} l_{ijk}) \quad (7)$$

The objective of the problem is thus placing δ detectors on the grid such that Equation (7) is minimized (this actually implies minimizing the second term therein since the first one is independent of the placement of the detectors). We shall denote this problem as the Optimal Placement of Suicide-Bomber Detectors (OPSBD). Notice that the OPSBD share some similarities with maximal location covering problems (Farahani et al, 2012), namely the fact that *facilities* (i.e. detectors) are to be located in order to service some *customers* (i.e. objectives). The OPSBD problem has its own distinctive features though. For example, the distance from detectors to objectives is to a large extent irrelevant (as long as it is larger than the neutralizing distance) and the interaction between detectors is highly non-linear. Alternatively, one could think of each cell through which a path goes as a “customer” (and again each of them would have a different contribution depending on how many paths go through it, the population density of the objectives those paths go to, the distance from the detector, and the non-linear interaction among the latter).

3 Description of Algorithms

In this section, we describe the different algorithms for the OPSBD problem that we have considered in this work. To this end, we shall start by presenting some general algorithmic considerations that are applicable to all algorithms considered, namely a cache data structure aimed to avoid recomputation and a dominance criterion that reduces the search space. Subsequently, we will provide a detailed description of all heuristics, namely a Greedy Randomized Adaptive Search Procedure (GRASP), a Hill Climber (HC), an Evolutionary Algorithm (EA) and a Univariate Marginal Distribution Algorithm (UMDA). Prior to these and for the sake of completeness, we will also describe the Greedy algorithm proposed by Nie et al (2007), which will be used to benchmark the remaining procedures.

3.1 General Algorithmic Considerations

All algorithms compared in this work use a cache data structure in order to accelerate the otherwise repetitive computations needed during the optimization of a particular instance.

This memory consists of several components. The first one is a three-dimensional array $\Lambda = \{\Lambda_{ijp}\}_{m \times n \times r}$ whose first two dimensions correspond to the number of rows (m) and columns (n) of the map. The third one corresponds to the total number of paths going from each entrance in the map to each objective ($r = \epsilon\phi$). We assume these are numbered according to some arbitrary predefined order, and use the notation i^p, j^p to denote respectively the entrance and objective corresponding to the p -th path under said ordering. Each entry in this data structure stores the length of the segment of the path that would be detected by placing a detector at the center of the cell denoted by the two first dimensions, i.e., if detector d_k is placed at cell (i', j') , $\Lambda_{i'j'p} = l_{i'j'pk}$, where $l_{i'j'pk}$ is as defined in previous section for the corresponding values i^p, j^p, k .




Fig. 4: Dominated cells in the map from Figure 2, assuming that a total of $\delta = 3$ detectors are to be placed. The detector effectiveness radius is $\tau = 10\text{m}$. Red bordered cells can be safely disregarded as candidate locations for placing the detectors. Lightest cells are more dominated than darker ones.

Another component of the cache is a bi-dimensional array $\Delta = \{\Delta_{ij}\}_{m \times n}$, which stores the number of times each cell in the map is dominated by remaining cells. We say that a cell $c_1 = (i_1, j_1)$ is dominated by another cell $c_2 = (i_2, j_2)$ if,

$$\forall p \in 1, \dots, \epsilon\phi : \Lambda_{i_1 j_1 p} \leq \Lambda_{i_2 j_2 p} \quad (8)$$

$$\exists p \in 1, \dots, \epsilon\phi : \Lambda_{i_1 j_1 p} < \Lambda_{i_2 j_2 p} \quad (9)$$

that is, for all paths in the map, the length of the segment of that path that would be detected by placing a detector at cell c_1 is less than or equal to the one detected by placing instead the detector at cell c_2 , and, for at least one path, the path detected by a detector at c_1 is strictly less than that detected from c_2 . This concept of dominance can be used to prune the search space of the problem since those cells that are dominated by a number of cells greater than or equal to the total number of detectors to be placed at the map (i.e., $\Delta_{ij} \geq \delta$) can be pruned, since there exists at least δ cells better than those. Figure 4 illustrates dominance values for the scenario shown before.

3.2 Greedy Algorithm

As mentioned above, this algorithm (denoted as Greedy henceforth) corresponds to the one originally described in (Nie et al, 2007) and will be used to establish a baseline for comparison purposes. This greedy heuristic is a constructive method that places the δ detectors one at a time as depicted in the pseudo-code provided in Algorithm 1.

Firstly, the set of candidate detectors is initialized with all non-blocked cells in the map minus those that are dominated by at least δ cells (line 1). The algorithm starts from an empty solution (line 2) and from there on, the δ detectors are successively

Algorithm 1: Greedy Algorithm

```

Input:  $\delta$  (number of detectors to be placed)
1  $candidates := \{(r, c) \mid 1 \leq r \leq m, 1 \leq c \leq n, A_{rc} \text{ is unblocked}, \Delta_{r,c} < \delta\}$ ;
2  $sol := \emptyset$ ;
3 while  $length(sol) < \delta$  do
4    $fitness^* := \infty$ ;
5    $detector^* := \text{null}$ ;
6   for  $d \in candidates$  do
7      $tentative := sol + \{d\}$ ;
8     if  $value(tentative) < fitness^*$  then
9        $detector^* := d$ ;
10       $fitness^* := value(tentative)$ ;
11    end
12  end
13   $sol := sol + \{detector^*\}$ ;
14   $candidates := candidates \setminus \{detector^*\}$ ;
15 end
16 return  $sol$ ;

```

added to the current partial solution sol (lines 3-15). In each iteration, the candidate detector whose incorporation leads to the best extended solution –i.e., locally minimizing Equation (7)– is selected (lines 6-12), added to current solution (line 13) and removed from the set of candidate detectors (line 14).

As it can be seen, this greedy algorithm performs different locally optimal decisions which can be taken with a low computational cost. However, due to the myopic functioning of this scheme, the resulting final solution is not guaranteed to be optimal. Although it was shown in (Nie et al, 2007) that the quality of solutions provided by the greedy algorithm for this particular problem is not very far from the global optimum for the relatively small-sized instances that were considered in that work, it will be later shown in Section 4 that other heuristics can provide better solutions for this problem if larger instances are considered.

3.3 Greedy Randomized Adaptive Search Procedure

The Greedy Randomized Adaptive Search Procedure (GRASP) (Feo and Resende, 1995; Festa and Resende, 2009a,b) is a metaheuristic that tries to alleviate the myopic functioning associated with greedy algorithms. This is also a constructive algorithm that adds one component of the solution at the time, but unlike the greedy algorithm, the one leading to largest local improvement is not always selected.

The pseudo-code in Algorithm 2 is a description of the specific incarnation of the general GRASP scheme that we have considered in this work. First of all, a *candidate list* (CL) of partial solutions is built by extending the current partial solution (sol) with all candidate detectors (lines 4-8). Subsequently, a *restricted candidate list* (RCL) is built by selecting some detectors from the candidate list (lines 9-17). This is done by using a threshold for the quality of tolerable partial solutions (μ), which we compute as a percentage of the difference between the values of the best and worse solutions in the candidate list (lines 9-11). Parameter α ($0 \leq \alpha \leq 1$) of the algorithm controls

Algorithm 2: GRASP Algorithm

```

Input:  $\delta$  (number of detectors)
          $\alpha$  (greedy temperation parameter)
1 candidates :=  $\{(r, c) \mid 1 \leq r \leq m, 1 \leq c \leq n, A_{rc} \text{ is unblocked}, \Delta_{rc} < \delta\}$ ;
2 sol :=  $\emptyset$ ;
3 while length(sol) <  $\delta$  do
4   CL :=  $\emptyset$ ;
5   for  $d \in \textit{candidates}$  do
6     tentative := sol +  $\{d\}$ ;
7     CL := CL  $\cup$   $\{\textit{tentative}, d\}$ ;
8   end
9   min :=  $\arg \min_{(s,d) \in \textit{CL}} \textit{value}(s)$ ;
10  max :=  $\arg \max_{(s,d) \in \textit{CL}} \textit{value}(s)$ ;
11   $\mu$  := min +  $\alpha(\textit{max} - \textit{min})$ ;
12  RCL :=  $\emptyset$ ;
13  for  $(s, d) \in \textit{CL}$  do
14    if  $\textit{value}(s) \leq \mu$  then
15      RCL := RCL  $\cup$   $\{(s, d)\}$ ;
16    end
17  end
18   $(\textit{sol}, d) := \text{RANDOMSELECT}(\textit{RCL})$ ;
19  candidates := candidates  $\setminus \{d\}$ ;
20 end
21 return sol;

```

the degree of greediness used. For instance, if α is 0, then only the best extended solution is included in the *RCL* (this would correspond to a greedy selection strategy). For greater values of α , more tentative decisions can be included in the *RCL* (all of them would be included when $\alpha = 1$) so that selection is less intense and the search is diversified. The next detector to be added to the solution is randomly chosen from the *RCL* in an uniform way (line 18). This process concludes when the solution being built is complete, i.e., it includes δ detectors. Although it is not shown in the pseudo-code, this whole process can be repeated until the maximum execution time is exhausted. The best solution found is then returned.

3.4 Hill Climbing

A simple and effective class of optimization heuristics based on the notion of *neighborhood* is Local Search (Aarts and Lenstra, 1997; Hoos and Stützle, 2005). As opposed to constructive methods, Local Search algorithms are procedures that work with complete solutions to the problem. In these algorithms, an initial solution is provided (which can either be constructed randomly or by means of another heuristic). Then, its neighborhood is examined, and if a better solution is found, a move is performed towards it. The process is repeated with the new solution until a local optimum is found.

Here, we consider a Hill Climbing (HC) algorithm. In the pseudo-code for this procedure (shown in Algorithm 3), *sol* denotes a complete solution to be improved,

Algorithm 3: Hill Climbing Algorithm

```

Input: sol (a collection with the coordinates of  $\delta$  detectors)
1 candidates :=  $\{(r, c) \mid 1 \leq r \leq m, 1 \leq c \leq n, A_{rc} \text{ is unblocked}, \Delta_{rc} < \delta\}$ ;
2 improvement := true;
3 while improvement do
4   improvement := false;
5   for i := 1 to length(sol) do
6     for d  $\in$  (candidates  $\setminus$   $\{d \mid d \in sol\}$ ) do
7       tentative := REPLACE(sol, i, d);
8       if value(tentative) < value(sol) then
9         sol := tentative;
10        improvement := true;
11      end
12    end
13  end
14 end
15 return sol;

```

represented as a collection of locations for different detectors. The algorithm substitutes one detector location in the current solution for all unused candidate locations for the problem instance (REPLACE(*sol*, *i*, *d*) denotes the solution that is obtained by replacing *i*-th detector in solution *sol* by alternative detector *d* – see line 7) and selects the one leading to a better solution (lines 8-11). If, as a result of this replacement, the new solution is better than the current one, the corresponding move is accepted (line 9) and an improvement is acknowledged (line 10). Using the new solution, the same process is repeated for the rest of detectors. The process is iterated as long as the solution has been improved, until no further enhancement is possible. The final solution is a local optimum and is returned as a result. Different starting solutions can be used and the whole scheme can be repeated until the maximum allowed execution time is reached.

3.5 Evolutionary Algorithm

Evolutionary algorithms (EAs) (Eiben and Smith, 2003) are black-box optimization procedures inspired by the biological evolution of species that work with a population of solutions subject to different operations such as reproduction, recombination, mutation and replacement. In this section, we describe a steady-state evolutionary algorithm that we have considered in order to tackle the OPSBD problem.

The pseudo-code of the EA is depicted in Algorithm 4. As shown, the algorithm uses a population (*pop*) of *popSize* non-repeated individuals, each one corresponding to one full solution to the problem. Each of these individuals is represented as a vector whose length is the number of detectors (δ) to be placed at the map. Hence, each element in the vector corresponds to one of the cells in the map where one of the detectors should be placed. Similarly to previous heuristics, the search space explored by the EA is restricted to non-blocked cells in the map minus those that are dominated by at least δ cells. Each solution in the population is initialized by randomly selecting

Algorithm 4: Evolutionary Algorithm

```

Input:  $popSize$  (population size)
          $p_X$  (recombination probability)
          $p_m$  (mutation probability)
1 for  $i := 1$  to  $popSize$  do
2    $pop_i := \text{RANDOMINDIVIDUAL}()$ ;
3    $\text{EVALUATE}(pop_i)$ ;
4 end
5 while allowed runtime not exceeded do
6   if recombination is performed( $p_X$ ) then
7      $parent_1 := \text{SELECT}(pop)$ ;
8      $parent_2 := \text{SELECT}(pop)$ ;
9      $offspring := \text{RECOMBINE}(parent_1, parent_2)$ ;
10  else
11     $offspring := \text{SELECT}(pop)$ ;
12  end
13   $offspring := \text{MUTATE}(p_m, offspring)$ ;
14   $\text{EVALUATE}(offspring)$ ;
15   $pop := \text{REPLACE}(pop, offspring)$ ;
16 end
17 return best solution found;

```

in an uniform way some of these cells (lines 1-4). Afterwards, the following process is repeated, until the maximum allowed execution time is reached:

- With probability of selection p_X , two parents are selected from the current population using *binary tournament selection* (lines 7-9). Let us now consider the set comprising the union of detector placements included in selected individuals. A new individual, constituting the offspring for this generation, is then defined by sampling in a random way δ different placements from this previous set.
- Otherwise (with probability $1 - p_X$) a random individual of current population is selected as the offspring (line 11).
- Each component of the offspring is mutated with probability p_m by replacing the corresponding detector in that component with another one not included in the solution (line 13), and the resulting individual is evaluated (line 14).
- As to replacement, the worst individual in the current population is replaced by the offspring (line 15).

Finally, when the execution time limit is reached, the best individual found is returned as a solution.

3.6 Univariate Marginal Distribution Algorithm

Finally, we have taken into account a Univariate Marginal Distribution Algorithm (UMDA) (Mühlenbein and Paaß, 1996), a metaheuristic which belongs to the class of Estimation of Distribution Algorithms (EDAs) (Larrañaga and Lozano, 2002). EDAs are themselves based on Evolutionary Algorithms, but whereas most classical EAs build new solutions using recombination and mutation operators, EDAs do this by

Algorithm 5: Estimation of Distribution Algorithm

```

Input: popSize (population size)
1 model := BUILDINITIALMODEL();
2 for i := 1 to popSize do
3   | popi := SAMPLEINDIVIDUAL(model);
4   | EVALUATE(popi);
5 end
6 while allowed runtime not exceeded do
7   | pop' := SELECT(pop);
8   | model := BUILDMODEL(pop');
9   | for i := 1 to popSize do
10  | | popi := SAMPLEINDIVIDUAL(model);
11  | | EVALUATE(popi);
12  | end
13 end
14 return best solution found;

```

sampling an explicit probability model which is updated along the optimization process. This model can be encoded in different ways, and, in the case of UMDAs, a simple univariate linear model is used. The general functioning of an EDA is illustrated in Algorithm 5.

In our case, we use a GRASP-based encoding (Cotta and Fernández, 2004), that is, individuals in the population encode the indexes in the RCL of the decisions that the GRASP algorithm would pick to place each detector, and the univariate probability distribution is used to model these indexes. More precisely, each solution is a vector $\mathbf{x} \in \mathbb{N}^{\delta-1}$, where x_i indicates that the i th detector is picked as the x_i th best (according to the greedy selection criterion) option available at that point, once that the previous $i - 1$ detectors have been placed (e.g., vector $\mathbf{x} = \mathbf{0}$ would encode the greedy solution described in Section 3.2 since all detectors would be picked according to the best local decision). Note that the length of \mathbf{x} is $\delta - 1$ because it does not make sense to place the last detector in any other location than the locally optimal one.

In the univariate model these variables are assumed to be independent and hence the probability distribution $p(\mathbf{x})$ is factorized as

$$p(\mathbf{x} = \langle v_1, \dots, v_{\delta-1} \rangle) = \prod_{i=1}^{\delta-1} p(x_i = v_i). \quad (10)$$

UMDA computes $p(x_i = v_i)$ as

$$p(x_i = v_i) = \frac{1}{|pop'|} \sum_{j=1}^{|pop'|} [pop'_{ji} = v_i], \quad (11)$$

where pop'_{ji} is the i th variable of the j th solution in pop' and $[\cdot] : \mathbb{B} \rightarrow \{0, 1\}$ is an indicator function ($[\text{true}] = 1$, $[\text{false}] = 0$).

At the beginning of the algorithm, the model is initialized (BUILDINITIALMODEL in line 1) in a heuristic way: the GRASP algorithm described in Section 3.3 is run $popSize$ times for a certain fixed value of α , and the decision vectors generated in

those runs are used to create the initial model as indicated in Equation 11. Function `SAMPLEINDIVIDUAL` is used to create the individuals in the initial population (line 3) by sampling the probabilistic model built. These individuals are decoded (using a guided GRASP algorithm as mentioned before) in order to be evaluated (line 4). From there on, the algorithm goes into an iterative process until the allowed execution time is exhausted. During each iteration, a new population of selected solutions from the current population is built by using truncation of the lower half (quality-wise) of the population (line 7), a new model is created by using the greedy order of detectors in these solutions (line 8), and a new population to be used in the next iteration is generated by sampling the new model (lines 9-12).

4 Experimental Results

In this section, we analyze the results of different experiments we have carried out regarding the OPSBD problem. First of all, we do a comparison of the performance of the various algorithms introduced in Section 3 on a extensive set of synthetic instances comprising a broad combination of parameters. Next, we do a sensitivity analysis on the performance of the most effective algorithm found in previous comparison, with the aim of understanding how different settings for parameters are related to the difficulty of problem instances. Finally, we evaluate how the best algorithm performs in the case of instances representing real world locations.

Although, as mentioned previously, we have considered a large set of settings for doing these experiments, there are some parameters for which we have used values which are good approximations to their typical configurations in a real environment and/or are analogous to the values previously used in the literature (Kaplan and Kress, 2005; Nie et al, 2007). In this way, we have assumed for simplicity that the probability of the attacker for choosing each of the paths going from one entrance to an objective is constant, i.e., $\gamma_j = 1/\phi\epsilon$, as in (Nie et al, 2007). We have also contemplated that effective neutralization of the attacker is not possible if the distance to the objective is less than 10m (here we assume that aborting the attack is plausible if the attacker is at least 10 seconds away from the objective and that the terrorist moves at a speed of 1m/s). In the case of the detector's instantaneous detection rate, we used a value of $\eta = 0.06$ and for the probability of effectively neutralizing a detected attacker, we used $\theta = 0.6$. Finally, in order to estimate the expected number of casualties for an objective cell C_j , $1 \leq j \leq \phi$, we used equation [2] in (Kaplan and Kress, 2005) which assumes that the number of fragments after the explosion tends to ∞ and that these and individuals around the target area follow a spatial Poisson process. We used the same parameters for this equation as those used in (Nie et al, 2007), except for the population densities near objective cells, which was set as a constant there but we modelled as a random variable in $\mathcal{N}(0.4, 0.1)$ persons / m² instead, with the aim of considering more diverse instances (see previous references for full details).

Regarding the hardware platform used, all experiments in this paper were executed on Intel Xeon E7-4870 2.4 GHz processors with 2GB of RAM running under SUSE Linux Enterprise Server 11 operating system.

4.1 Random instances

In this section we present results of an experimental comparison on the performance of different proposed heuristics for the OPSBD problem. We have considered the following algorithms with the parameterization indicated:

- The local search Hill Climbing algorithm presented in Section 3.4. This algorithm was run on different random generated instances until the allowed execution time was exhausted.
- The Evolutionary Algorithm described in Section 3.5. For this algorithm, we used standard parameters and operators: a population size $popSize = 100$ individuals, probability of crossover $p_X = 0.9$, probability of mutation $p_m = 1/\delta$, binary tournament selection mechanism and replacement of the worst individual.
- The Greedy Randomized Adaptive Procedure introduced in Section 3.3. In this case, we have considered different settings for parameter $\alpha \in \{0.1, 0.25, 0.5\}$. We denote these by $GRASP_{\alpha=0.1}$, $GRASP_{\alpha=0.25}$ and $GRASP_{\alpha=0.5}$ respectively to indicate the value of the parameter in each case. In addition, we have considered a version of this same heuristic that performs a local search on each constructed solution by using the Hill Climbing algorithm. We denote this latter variant by $GRASP_{\alpha=0.1}+HC$ (we have considered this hybrid approach just for $\alpha = 0.1$).
- The Greedy algorithm originally introduced in (Nie et al, 2007).
- The Univariate Marginal Distribution Algorithm described in Section 3.6. The parameters for this algorithm were: population size $popSize = 100$ individuals and a selected population pop' of 50 individuals. Three different settings for parameter α were used for building the initial probabilistic model, leading to three versions of the algorithm: $UMDA_{\alpha=0.1}$, $UMDA_{\alpha=0.25}$ and $UMDA_{\alpha=0.5}$.

All algorithms have been run on random instances constructed using the following combinations of parameters:

- Maps with dimensions of 32×32 and 64×64 cells.
- Different side lengths for cells comprising the maps: $\zeta \in \{5, 10, 20\}$ meters.
- Different number on entrances on each border of the map in $\{2, 3, 4\}$ leading to $\varepsilon \in \{8, 12, 16\}$.
- Different number of objective cells on the map: $\phi \in \{2, 4, 6, 8\}$.
- A percentage of blocked cells on the map of 5%.
- Different number of detectors to be placed on the map: $\delta \in \{6, 8, 10\}$.
- A detection radius $\tau = 20$ meters.

For each combination of these parameters, 25 random instances were considered, thus this benchmark comprises a total of 5,400 different problem instances². Following (De Jong et al, 1997), each algorithm is run once on each instance thus allowing to obtain a more representative measure of performance across the whole set of possible instances than that obtained by performing multiple runs on a reduced hand-picked set of instances for which a certain algorithm could have been better adapted. A maximum execution time of 5 seconds was allowed for each algorithm in the case of 32×32 instances, whereas the execution time allowed for 64×64 instances was 10

² These instances are publicly available – see the Appendix for details.

Table 1: Mean ranks of all algorithms on 32×32 and 64×64 cells instances.


position	algorithm	ranking
1 st	HC	1.28
2 nd	GRASP $_{\alpha=0.1}$ +HC	1.83
3 rd	EA	3.17
4 th	GRASP $_{\alpha=0.1}$	4.82
5 th	UMDA $_{\alpha=0.1}$	4.93
6 th	GRASP $_{\alpha=0.25}$	6.08
7 th	UMDA $_{\alpha=0.25}$	6.75
8 th	Greedy	7.39
9 th	GRASP $_{\alpha=0.5}$	8.95
10 th	UMDA $_{\alpha=0.5}$	9.80

seconds. Since the cache data structure described in Section 3.1 has to be precomputed by all compared algorithms, its computation time is not included in these time limits.


In order to measure performance in a comparable way across the huge collection of problem instances considered, we have recorded the best results found by any algorithm on each instance, and used these to calculate relative deviations to such best known solutions. The results are summarized in Figure 5 showing box plots of the distribution of relative deviations to best known solutions for different algorithms, separately for the 32×32 and 64×64 instances. It can be seen that deviations are very close to 0 for HC, which means that this algorithm was able to provide best known solutions for most of the instances. The EA comes third, with distances that are also quite close to 0. GRASP $_{\alpha=0.1}$ and GRASP $_{\alpha=0.25}$ are the next two better algorithms, but their performance degrades clearly when compared to HC and EA. The value of upper whisker for the distribution corresponding to Greedy on 32×32 instances is of a 4.59% relative distance to the optimum, whereas the same value for HC is 0% and for EA is 0.48%. The difference is of a 5.78% in the case of 64×64 instances. This is an indication of the extent of improvements achieved by both heuristics over Greedy.

Subsequently, a statistical test was used in order to analyze whether differences among the different techniques are statistically significant. More precisely, the Friedman Test (FT) (Friedman, 1937, 1940), a non-parametric test based on rankings, has been used for this purpose. If, as a result of this test, the null hypothesis stating equality of rankings between the different techniques is rejected, we proceed to post-hoc procedures in order to detect statistical differences among different pairs of algorithms. For this purpose, Shaffer’s static procedure (Shaffer, 1986) with a standard significance level of $\alpha = 0.05$ has been used. All these analysis were carried out using the software package provided by the Soft Computing and Intelligent Information Systems group at University of Granada (García and Herrera, 2008).

Firstly, the mean rank of all ten algorithms in the benchmark considered is computed and shown in Table 1. As can be seen, the order of performance from best to worst is: HC, followed by EA, the two variants of GRASP|UMDA using smallest values of α , the Greedy algorithm and finally GRASP|UMDA $_{\alpha=0.5}$. To determine



(a)



(b)

Fig. 5: Relative distances to best solutions of results provided by different algorithms on (a) 32×32 instances and (b) 64×64 instances.

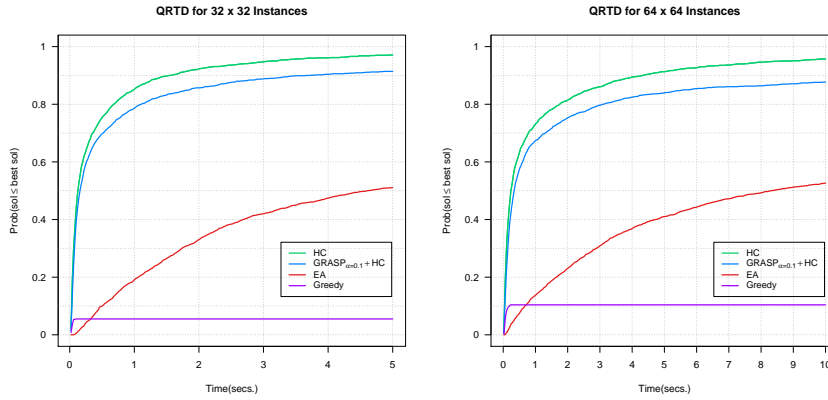


Fig. 6: Qualified runtime distributions (QRTD) for best performing algorithms (Greedy also included for comparison purposes) on 32×32 (left) and 64×64 (right) instances. Curves show the probability of reaching best known solutions along time for different algorithms.

the significance of these rank differences, the FT is performed and results in $\chi_F^2 = 1726.18$, much larger than the critical value 16.92 for $\alpha = 0.05$ (according to χ^2 with 9 degrees of freedom). The p -value calculated by this test is actually $p = 0$ which provides strong evidence for rejecting equality of rankings. Shaffer's post-hoc procedure, at the same standard $\alpha = 0.05$ level, found statistically significant differences between all pairs of algorithms, except for $\text{GRASP}_{\alpha=0.1}$ vs $\text{UMDA}_{\alpha=0.1}$, $\text{GRASP}_{\alpha=0.25}$ vs $\text{UMDA}_{\alpha=0.25}$, $\text{GRASP}_{\alpha=0.5}$ vs $\text{UMDA}_{\alpha=0.5}$, $\text{UMDA}_{\alpha=0.25}$ vs Greedy, and HC vs $\text{GRASP}_{\alpha=0.1} + \text{HC}$. It can thus be concluded that using the $\text{GRASP}_{\alpha=0.1}$ component for initializing solutions for the HC algorithm is not advantageous with respect to using random initialization for this same purpose. Moreover, using the UMDA in order to guide the search of GRASP does not provide benefits with respect to taking random decisions, as it is done in the pure GRASP algorithm. From a more general point of view, all the algorithms proposed in this paper have a better performance than the Greedy algorithm from (Nie et al, 2007), except $\text{GRASP}_{\alpha=0.5}$ and $\text{UMDA}_{\alpha=0.5}$ (statistically worse) and $\text{UMDA}_{\alpha=0.25}$ (better than Greedy but not significantly so). It can be also seen that the performance of GRASP algorithm improves with smaller values of α , but this algorithm performs worse than HC and EA for this problem.

The behavior as anytime algorithms of both the best performing heuristics and Greedy is analyzed in Figure 6. This figure shows Qualified Runtime Distributions (QRTD) for the best three algorithms, namely HC, $\text{GRASP}_{\alpha=0.1} + \text{HC}$ and EA on 32×32 and 64×64 instances. Such QRTDs show the probabilities of reaching best known solutions as a function of execution time. It can be seen that the algorithm that progresses faster towards high quality solutions is HC. This algorithm is able to find best solutions for 90% of the 32×32 instances in about 1.5 seconds. For 64×64 instances, the same rate of success is achieved after 4.0 seconds of execution. Tak-




Fig. 7: Relative distances to best known solutions of results provided by HC and Greedy on 128×128 instances.

ing into account the context in which the OPSBD problem should be solved, these seem to be sensible time requirements, amenable for fast-response adjustment of detector locations. From a more detailed point of view, the previous analysis showed that HC and $\text{GRASP}_{\alpha=0.1}+\text{HC}$ were roughly equivalent with respect to the quality of solutions they provide, but QRTDs show that $\text{GRASP}_{\alpha=0.1}+\text{HC}$ is however slightly slower than HC. As $\text{GRASP}_{\alpha=0.1}+\text{HC}$ is a more complicated algorithm to implement and also does not provide advantages over HC, we can conclude that using the latter one is a better alternative to solve this problem. Regarding the EA, QRTDs also show that it progresses at a slower rate than HC, as its curve is always below the one for HC. Finally, although Greedy is not able to progress towards best solutions (it is a deterministic constructive algorithm) its inclusion in the QRTDs allows comparing its performance to the remaining algorithms on a time-quality basis. Notice that compared to the EA, it provides a better solution initially, but the EA can outperform Greedy in a matter of tenths of a second. This is exactly the expected behavior for such a kind of algorithm, as Greedy local decisions can be made in a very fast way but they are expected to be suboptimal in general.

Finally, we compared the best performing algorithm (HC) with the previous proposal in the literature (Greedy) on larger random instances. In this case, maps consisted of 128×128 grids of cells and, for the rest of parameters, we considered the same combinations as those we used previously for 32×32 and 64×64 instances. Maximum allowed execution times were 10 seconds per instance in this case. A Wilcoxon signed ranked statistical test rejects the null hypothesis stating equality

of rankings between the two algorithms (p -value=0). Boxplots for relative distances to best known results are shown in Figure 7. It can be seen that HC algorithm also provides consistently best results for these instances, and that Greedy is not able to provide solutions with same quality. The value of upper whisker for the distribution corresponding to these instances of Greedy is at a 5% relative distance to the optimum, whereas the same value for HC algorithm is 0%.

Notice that the number of cells in the instance mostly affects the size of the search space and the ability for fine-tuning solutions. Results of our experimentation indicate that this parameter does not exert a strong influence on the quality of the results though, as the value of upper whisker for the distribution corresponding to the relative distances between the solutions provided by Greedy and HC is consistently close to 5% for the different settings of this parameter that we have considered: 32x32, 64x64 and 128x128.

4.2 Sensitivity Analysis

In this section, we perform a sensitivity analysis on the performance of the Hill Climbing algorithm (Section 3.4) with respect to Greedy (Section 3.2) with the aim of studying which parameter settings of the problem lead to a larger performance gap among these two techniques. The rationale for this procedure is that if the relative improvement of the solution provided by Hill Climbing over Greedy increases for some parameter settings, it can be considered that those instances turn out to be harder-to-solve for the latter. The different parameters and settings that have been analyzed are a superset of those considered in the previous section, namely:

- Maps with dimensions of 32×32 and 64×64 cells.
- Different side lengths for cells comprising the maps: $\zeta \in \{5, 10, 20\}$ meters.
- Different number on entrances on each border of the map in $\{2, 3, 4\}$ leading to $\varepsilon \in \{8, 12, 16\}$.
- Different number of objective cells on the map: $\phi \in \{2, 4, 6, 8\}$.
- Different percentages of blocked cells on the map: $\varpi \in \{2.5\%, 5\%, 10\%\}$.
- Different number of detectors to be placed on the map: $\delta \in \{2, 4, 6, 8, 10\}$.
- Different values of the detection radius: $\tau \in \{10, 20, 40\}$ meters.

All possible combinations of these parameter settings lead to 3,240 different kinds of instances, and for each of them, we have considered 25 random instances, so that our analysis comprises a total of 81,000 different instances. In the case of 32×32 cell maps, the maximum execution time allowed for the algorithms on each instances was 5 seconds whereas in the case of 64×64 maps, this time was extended to 10 seconds.

4.2.1 Percentage of Blocked Cells

We firstly analyzed differences between maps with varying percentages of blocked cells ($\varpi \in \{2.5\%, 5\%, 10\%\}$). The ranking is shown in Figure 8. The value of the Friedman statistic is $\chi_F^2 = 7.59$ that corresponds to a p -value = 0.02 (the critical value in this case is distributed according to χ^2 with 2 degrees of freedom, yielding




Fig. 8: Relative improvements of HC over Greedy for different percentages of blocked cells (ϖ) in the map. Settings are ordered (left to right) from more difficult to easier ones. Rankings are shown at the right side of labels. In order to interpret this figure and subsequent ones in this subsection, notice that the boxplots depict the distribution of the aggregated data and show the range of variation in each case, whereas the rankings are computed by aligning the results obtained by each algorithm on the instances generated by each parameter combination and are used by the statistical tests.

5.99). This provides evidence for rejecting the null hypothesis that states equality of rankings between the different settings for ϖ . The order in Figure 8 goes from more difficult instances to easier ones, so that the first position corresponds to the setting for which the Hill Climbing algorithm provides the largest improvement. Post-hoc procedures only show statistical significant differences, at the standard significance level of $\alpha = 0.05$, between configurations for ϖ parameter of 2.5% vs 10%. Thus, it can be seen that a lower percentage of blocked cells leads to maps that are harder to solve for Greedy, but also that there are no differences between the intermediate setting for this parameter and the smaller or larger ones. We interpret this result in terms of the broader dispersion of shortest paths in instances with fewer obstacles, which makes more difficult to find major cross points and hence the goodness of solutions relies on the distribution of detectors with a more global perspective, as opposed to covering a couple of critical points and use the remaining detectors for fine tuning.




Fig. 9: Relative improvements of HC over Greedy for different side lengths (ζ) of cells in the grid. Settings are ordered (left to right) from more difficult to easier ones. Rankings are shown at the right side of labels.

4.2.2 Side Length of Cells

Regarding side lengths of cells (ζ) in the grid, the ranking is shown in Figure 9. We obtain $\chi_F^2 = 518.25$, corresponding to a p -value = $2.34 \cdot 10^{-10}$ (the critical value is here distributed again according to χ^2 with 2 degrees of freedom, yielding 5.99). Thus, there is strong evidence for rejecting the null hypothesis that states equality of rankings between the different settings for ζ . In this case, post-hoc procedures found statistically significant differences (at the standard significance level of $\alpha = 0.05$) between all possible settings for this parameter. It can be seen that the difficulty of instances increases as cell sizes decrease. One possible explanation for this phenomenon may be that having a more fine-grained map allows for a more precise placement of detectors in the map. This seems to be better exploited by Hill Climbing, that has more choices for fine-tuning the solution from a global perspective (as opposed to the local approach of Greedy).

4.2.3 Detection Radius

The detectors placed in the map may have a different radius of effectiveness. Here we consider settings for this parameter in $\tau \in \{10m, 20m, 40m\}$. The resulting ranking is shown in Figure 10. We obtain $\chi_F^2 = 825.17$, corresponding to a p -value = $2.9 \cdot 10^{-10}$ (the critical value is once again distributed according to χ^2 with 2 degrees

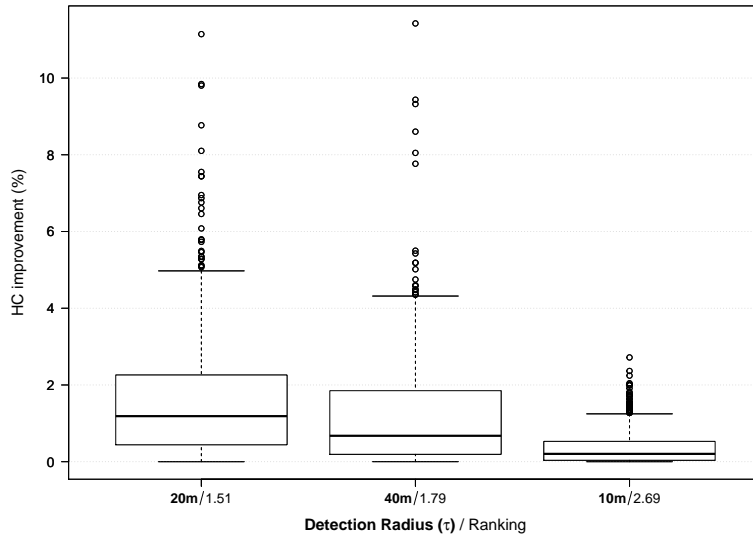


Fig. 10: Relative improvements of HC over Greedy for different values of the detection radius (τ). Settings are ordered (left to right) from more difficult to easier ones. Rankings are shown at the right side of labels.

of freedom). Hence, the null hypothesis stating equality of rankings between the different settings for this parameter should be rejected. In the case of this parameter, post-hoc procedures found statistically significant differences (at the standard significance level of $\alpha = 0.05$) between all pairs of settings. It must be noted that for this parameter harder instances are those with an intermediate radius of detection. One interpretation of these results may be that a smaller radius leads to less interactions between the detectors and paths followed by terrorists, whereas a very large radius allows for one detector to cover a great extent of the map, thus simplifying the difficulty of the problem. It is the intermediate radius $\tau = 20\text{m}$ that leads to harder instances in this case as there is a good degree of interaction of each detector with different paths and, at the same time, the radius is small enough for a precise placement of the detector to turn out to be crucial. Of course, this effect also depends on the number of detectors to be placed since, if this number is large, a suboptimal placement of some detectors may be to some extent compensated by other detectors covering the same paths. In this case, the smaller the detection radius, the more influential the location of the detectors would become.

4.2.4 Number of Detectors

We have considered different number of detectors (δ) to be placed on the map. Rankings are shown in Figure 11. The associated Friedman statistic is $\chi_F^2 = 763.50$ which




Fig. 11: Relative improvements of HC over Greedy for different number of detectors (δ) placed on the map. Settings are ordered (left to right) from more difficult to easier ones. Rankings are shown at the right side of labels.

implies a p -value = $2.41 \cdot 10^{-10}$ (the critical value is in this case distributed according to χ^2 with 4 degrees of freedom, yielding 9.49) and thus there are significant differences for different settings of this parameter. Post-hoc procedures found statistically significant differences (at the standard significance level of $\alpha = 0.05$) between all pairs of settings except for 6 vs. 8. Results show that instances with a smaller number of detectors ($\delta \in \{2, 4\}$) are the easiest ones, but also that instances with the maximum number of detectors ($\delta=10$) are easier than those with 8 and 6 detectors. One interpretation of these results is that the search space for a small number of detectors is relatively easy to explore for different algorithms, but also that having a higher number of detectors allows for placing them in a less precise way without affecting in a significant way the quality of the solution. In other words, instances with an intermediate number of detectors lead to a large enough search space which raises the difficulty of the problem. At the same time, a precise placement of the limited number of available detectors is required in order to provide high-quality solutions for these instances.

4.2.5 Number of Entrances

Regarding the number of entrances on the map (ϵ), we have considered 2, 3 and 4 entrances on each side of the map, which leads to the following settings: $\epsilon \in \{8, 12, 16\}$. Rankings are shown in Figure 12. The Friedman statistic is $\chi_F^2 = 202.47$ and the cor-

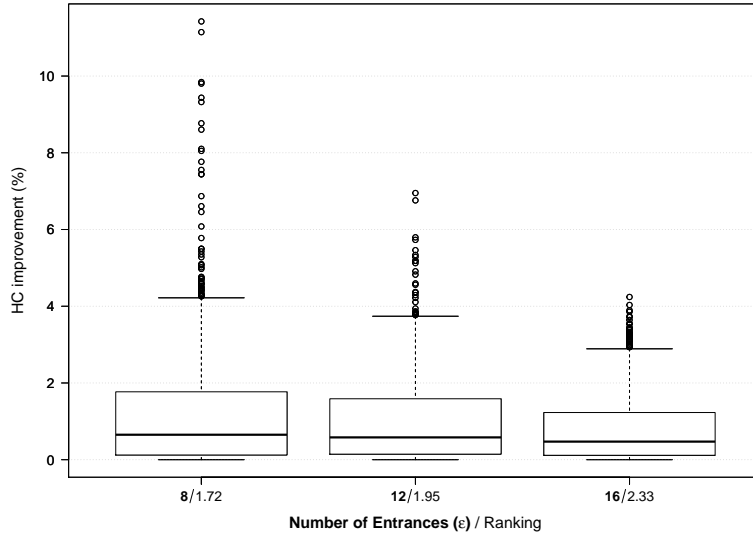


Fig. 12: Relative improvements of HC over Greedy for different number of entrances (ϵ) on the map. Settings are ordered (left to right) from more difficult to easier ones. Rankings are shown at the right side of labels.

responding p -value is $9.97 \cdot 10^{-11}$ (the critical value is in this case distributed according to χ^2 with 2 degrees of freedom, yielding 5.99). Post-hoc procedures found differences (at the standard significance level of $\alpha = 0.05$) between all pairs of settings for this parameter. In this case, the larger the number of entrances, the easier the resulting problem instances are. This result can be interpreted in terms of the number of paths (which is proportional to the number of entrances). Recall that the probability of targeting a certain objective is distributed across the different paths emanating from each entrance. Therefore, the larger the number of entrances, the smaller the weight of individual paths. This means that adjusting the precise location of a detector will yield a smaller variation of the objective function as a result of the the different coverage of the numerous existing paths. Of course, we cannot take this argument to the opposite limit because if the number of entrances was very low it would be easy to cover the paths by any appropriate heuristic.

4.2.6 Number of Objectives

Here we consider maps with different numbers of objective cells ($\phi \in \{2, 4, 6, 8\}$). Rankings are shown in Figure 13. The Friedman statistic is $\chi_F^2 = 94.69$, implying a p -value = $5,74 \cdot 10^{-11}$ (the critical value is in this case distributed according to χ^2 with 3 degrees of freedom, yielding 7.81). Post-hoc procedures in this case show that there are not statistically significant differences (at the standard significance level of

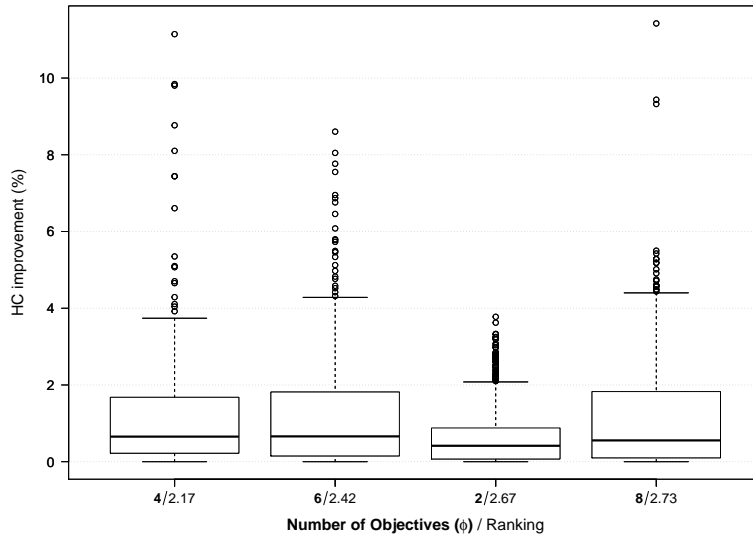


Fig. 13: Relative improvements of HC over Greedy for different number of objectives (ϕ) on the map. Settings are ordered (left to right) from more difficult to easier ones. Rankings are shown at the right side of labels.

$\alpha = 0.05$) between maps with 2 and 8 objectives. In this case, instances with the lowest ($\phi = 2$) and highest ($\phi = 8$) number of objectives are thus indistinguishable, and both lead to easiest-to-solve instances. For intermediate settings of this parameter, instances with $\phi = 4$ are harder-to-solve than those with $\phi = 6$. These results are partly explained in accordance to the argument laid down in Section 4.2.5 for the number of entrances, since the number of paths is also proportional to the number of objectives. In this sense, the lower limit ($\phi = 2$) correspond to instances in which there can be little variation of quality in solutions because most paths leading to objectives can be adequately covered. Similarly, in the upper limit ($\phi = 8$) the detectors need to cover much ground and as a result adjusting the location of a detector placed by Greedy can increase the coverage of some paths at the expense of others, resulting in the smallest window of improvements.

4.3 Real Instances

Finally, we have deployed Hill Climbing and Greedy on some more realistic instances resembling the features exhibited by urban areas that might be subject to a terrorist attack. These are depicted in Figures 14a–14c and comprise three different scenarios: (1) a large square, (2) an old-town area and (3) a mixture of the previous two – see Table 2. All of them have been constructed based on real-world locations. Objec-

Table 2: Description of the three realistic scenarios considered in the experimentation.

Name	Area (m ²)	Number of entrances (ϵ)	Number of objectives (ϕ)	Description
map₁	160 × 160	22	61	A large square surrounded by buildings and with only a few entrances, in which several monuments and other obstacles are scattered.
map₂	224 × 224	16	25	An old-town area featuring small plazas and curvy, non-perpendicular streets that connect them.
map₃	352 × 352	11	9	A mixture of the two previous scenarios that includes a large square and some adjacent streets

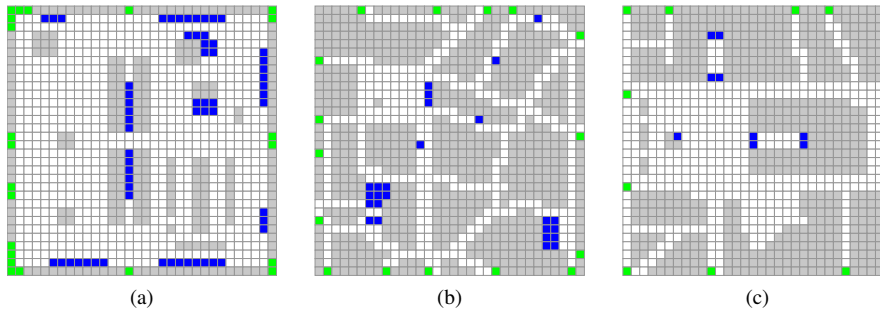


Fig. 14: Realistic maps considered. (a) A large plaza in the city center. (b) An old-town urban area with curvy streets and small plazas. (c) A mixed scenario including both a large plaza and some adjacent streets.

tives in these scenarios correspond to typical crowded places such as cafe terraces, small flea-markets or bazaars, or monuments with great touristic attractive and other iconic landmarks that gather people around them. All scenarios are represented on a 32×32 grid. These instances will be used to test the comparative performance of both algorithms (Hill Climbing and Greedy) in a more pragmatic context, identifying the circumstances under which there is a larger quality gap between them and how different the solutions provided by either algorithm are from a structural point of view.

The experimentation has been done with a number of detectors δ ranging from 1 to 50 and for different values of the detector radius $\tau \in \{10, 20, 40\}$ meters. All experiments have been replicated 25 times for each combination of map, number of detectors and detection radius. We let the Hill Climbing algorithm run for 10 seconds or until a local optimum was found (cf. Algorithm 3). This resulted in Hill Climbing taking an average of 4.45 seconds to provide its final solution (standard error $\sigma_{\bar{x}} = 0.043s$). Hill Climbing has consistently found the same solution in each case (Greedy is deterministic and therefore only one run is required per test case). The results are summarized in Figure 15a, in which the average relative improvement of

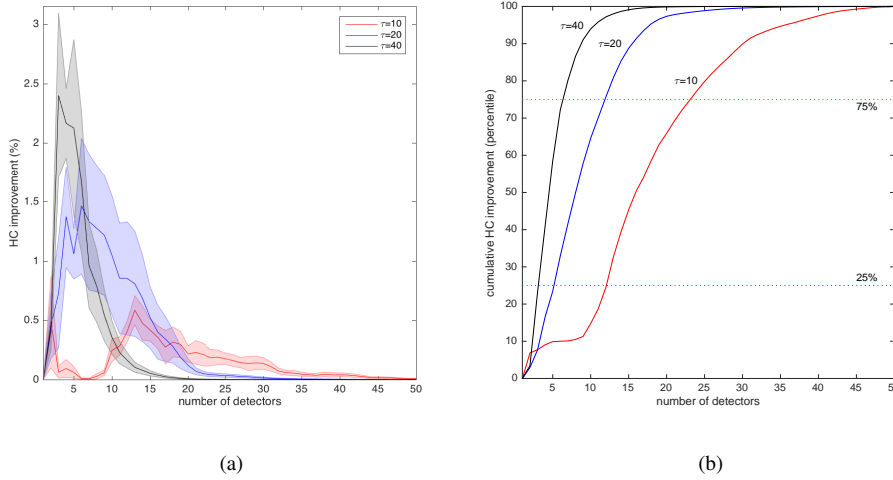


Fig. 15: (a) Relative improvement of solutions provided by HC over those provided by Greedy algorithm across the three scenarios, as a function of the number of detectors δ for different values of the detection radius τ . The shaded areas indicate the standard error. (b) Cumulative distribution of improvements depending on the detection radius τ . The inner quartiles are shown.

Hill Climbing over Greedy is shown. As can be seen, this improvement is always positive for any number of detectors. Indeed, this improvement is statistically significant (p -value ≈ 0) according to a Wilcoxon signed rank test for any scenario and detection radius, except for **map₃** and $\tau = 40$. Aggregating the results for all values of τ , the improvement is also statistically significant for all scenarios (again p -value ≈ 0). Not surprisingly, the relative improvement initially increases when more detectors are used, and then declines after reaching a certain peak. Obviously, as more and more detectors are placed the fitness function starts to saturate (recall Equation 7 – when the coverage of the paths is large, an additional increase in this coverage only implies a small fitness increase due to the exponential decay of the last term of the equation). It is nevertheless interesting to note two qualitative differences in behavior depending on the value of the detection radius. Firstly, as τ decreases, the peak is located in an interval comprising more detectors. This is more precisely depicted in Figure 15b, in which the cumulative distribution of improvement is shown (normalized in terms of percentiles for comparison purposes). If we focus on the central half of the distribution (that is, the second and third quartiles), it comprises $\delta \in [4, 7]$ for $\tau = 40$ m, $\delta \in [6, 12]$ for $\tau = 20$ m, and $\delta \in [13, 24]$ for $\tau = 10$ m. This can be interpreted in terms of the area covered by each detector in that case: as the detection radius increases, larger segments of the paths can be covered and therefore the fitness functions starts to saturate at an earlier point than it would for smaller detection radius. In the latter

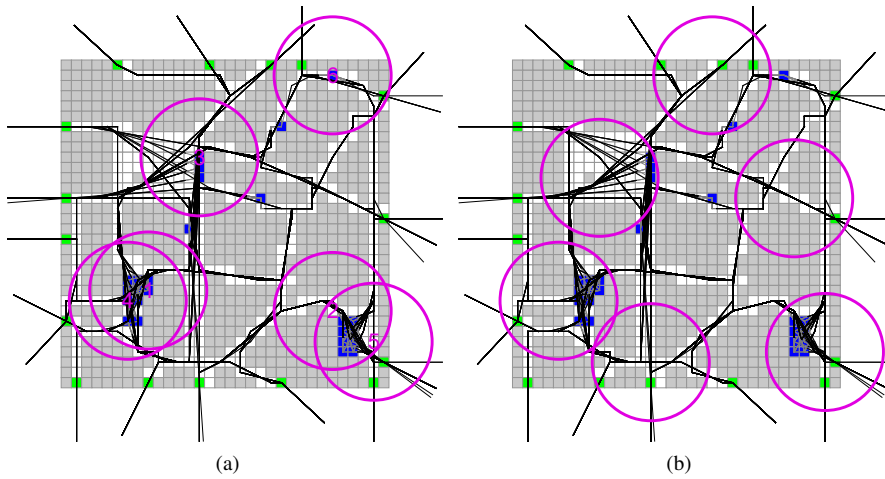


Fig. 16: Solutions found for **map₂** using $\delta = 6$ detectors of radius $\tau = 40\text{m}$. (a) Greedy (b) Hill Climbing. The numbers close to detector locations in the left image indicate the order in which Greedy places the detectors.

case, adjusting the location of detectors result in comparatively smaller length variations in the paths covered, implying that the fitness function saturates more slowly (and hence for a larger interval of detectors) and relative improvements are comparatively smaller than for larger values of the detection radius. Indeed, if we perform a head-to-head comparison among the values obtained on each map and number of detectors for each value of τ , we observe that $\tau = 10$ ranks better than $\tau \in \{20, 40\}$ (mean ranks are 1.48, 1.87 and 2.65 respectively), the differences being statistically significant (Friedman test p -value ≈ 0 , Shaffer's test passed at $\alpha = 0.05$ for all pairs of values of τ). This is in agreement with the interpretation laid down in Section 4.2 for random instances, and underlines the larger interval for which improvements are obtained for decreasing values of the detector radius.

Figure 16 show a comparison of the solutions found by Greedy and Hill Climbing on a specific case, namely the old-town scenario with $\delta = 6$ and $\tau = 40\text{m}$. Notice how the solution created by Greedy places the first three detectors trying to cover the densest target areas. Then, it tends to place the next two detectors in partial overlap with the first two, aiming to maximize the coverage of these areas by virtue of the independent functioning of the detectors. Finally, the last detector is placed close to a corner of the map, covering paths originating in two entrances. On the other hand, the solution provided by Hill Climbing has three detectors in locations rather close, but not identical, to the first three detectors placed by Greedy. This slight variation in their locations allows to place the remaining three in strategic positions to cover all paths originating in the four upper entrances, in three of the lower entrances, and in one of the entrances in the right part whose paths were receiving less attention by Greedy. Overall, the solution found by Hill Climbing has a more global rationale and looks more spatially balanced.

5 Conclusions

Suicide bombing is an infamous form of terrorism that presently causes a large number of casualties worldwide. In this work we have researched the problem of placing a number of non-fully reliable detectors on a threat area with known targets with the aim of maximizing chances of detecting a suicide-bombing attack, thus minimizing the expected number of casualties.

Apart from a branch-and-bound algorithm that does not scale well with problem size, the only previous proposal in the literature to tackle this problem was a greedy algorithm. We have approached this problem for the first time to the best of our knowledge from the point of view of iterative heuristics and metaheuristic techniques. To this end, we firstly performed an experimental comparison of the different proposed algorithms on a benchmark comprising a large set of random instances with different parameterizations. Results indicate that, among all considered techniques, a Hill Climbing algorithm obtains the best results in a consistent way, clearly outperforming the previous proposal from the literature, as well as other metaheuristic approaches. While the latter may be further improved by including other algorithmic components, it is interesting to note how such a relatively simple heuristic can provide a comparatively good performance, suggesting the appropriateness of search intensification in this context. Secondly, we did a sensitivity analysis in order to determine those settings of problem parameters leading to harder-to-solve instances. Thirdly, the Hill Climbing heuristic (as the best performing algorithm) was experimentally compared to the greedy algorithm (as the previous approach in the literature) on a set of real world scenarios that could be subject to terrorist attacks. Results of these tests corroborate the good performance of Hill Climbing observed before.

Different lines of research are open as future work to the present study. First of all, alternative algorithmic approaches can be explored. For instance, we have observed that the proposed UMDA is able to outperform GRASP in case the maximum allowed execution time is increased. Designing an alternative probabilistic model or performing new experimental comparisons with increased allowed execution times constitute promising extensions to this paper. Another promising proposal is to hybridize in a synergistic way the EA and Hill Climbing algorithms, as best performing algorithms studied in this work. To this end, Memetic Algorithms (Neri et al, 2012) constitute a well established and successful framework that can be used for this purpose. Also, more sophisticated versions of the problem studied could be also tackled in future work, like, for instance, considering more precise models of the functioning of detectors or making dynamic the locations of targets, so that their situation and even their number may vary along time. Notice finally that the interesting connection between the OPSBD problem and maximal covering location problems mentioned in Sect. 2 could be also used as a source of inspiration for designing new, more complex algorithmic approaches to this problem in future work.

Acknowledgements We would like to thank Mr. Antonio Hernández Bimbela for his help during the initial stage of this project, and to the anonymous reviewers for useful comments.

A Data Instances

The data instances used in the experimentation can be downloaded from:

- Synthetic instances:
 - 32×32 : <https://doi.org/10.6084/m9.figshare.4542691.v1>
 - 64×64 : <https://doi.org/10.6084/m9.figshare.4542700.v1>
 - 128×128 : <https://doi.org/10.6084/m9.figshare.4542706.v1>
- Realistic instances: <https://doi.org/10.6084/m9.figshare.4542817.v1>

References

- Aarts EHL, Lenstra JK (1997) Local Search in Combinatorial Optimization. John Wiley & Sons
- AFP (2016) Death toll from Baghdad blast rises to 292: minister. *France 24*. <http://www.france24.com/en/20160707-death-toll-baghdad-blast-rises-292-minister>, [Accessed 15 July 2016]
- Archavaleta G, Laumond JP, Hicheur H, Berthoz A (2008) An optimality principle governing human walking. *IEEE Transactions on Robotics* 24(1):5–14
- Carson JV (2017) Assessing the effectiveness of high-profile targeted killings in the “war on terror”. *Criminology & Public Policy* 16(1):191–220
- Caygill JS, Davis F, Higson SP (2012) Current trends in explosive detection techniques. *Talanta* 88:14–29
- Chicago Project on Security and Terrorism (CPOST) (2016) Suicide attack database (April 19, 2016 release). [Data File]. Retrieved from <http://cpostdata.uchicago.edu/>
- Cotta C, Fernández AJ (2004) A hybrid GRASP - evolutionary algorithm approach to golomb ruler search. In: Yao X, et al (eds) *Parallel Problem Solving from Nature - PPSN VIII*, Springer, Lecture Notes in Computer Science, vol 3242, pp 481–490
- De Jong KA, Potter MA, Spears WM (1997) Using problem generators to explore the effects of epistasis. In: Bäck T (ed) *Proceedings of the 7th International Conference on Genetic Algorithms*, Morgan Kaufmann, pp 338–345
- Edwards D, McMenemy L, Stapley S, Patel H, Clasper J (2016) 40 years of terrorist bombings – a meta-analysis of the casualty and injury profile. *Injury* 47(3):646–652
- Eiben AE, Smith JE (2003) *Introduction to Evolutionary Computing*. SpringerVerlag
- Farahani RZ, Asgari N, Heidari N, Hosseininia M, Goh M (2012) Covering problems in facility location: A review. *Computers & Industrial Engineering* 62(1):368 – 407
- Feo TA, Resende MGC (1995) Greedy randomized adaptive search procedures. *Journal of Global Optimization* 6(2):109–133
- Festa P, Resende MGC (2009a) An annotated bibliography of GRASP – Part I: Algorithms. *International Transactions in Operational Research* 16(1):1–24
- Festa P, Resende MGC (2009b) An annotated bibliography of GRASP – Part II: Applications. *International Transactions in Operational Research* 16(2):131–172
- Friedman M (1937) The Use of Ranks to Avoid the Assumption of Normality Implicit in the Analysis of Variance. *Journal of the American Statistical Association* 32(200):675–701
- Friedman M (1940) A comparison of alternative tests of significance for the problem of m rankings. *Ann Math Statist* 11(1):86–92
- García S, Herrera F (2008) An extension on “statistical comparisons of classifiers over multiple data sets” for all pairwise comparisons. *Journal of Machine Learning Research* 9:2677–2694
- Gares KL, Hufziger KT, Bykov SV, Asher SA (2016) Review of explosive detection methodologies and the emergence of standoff deep UV resonance raman. *Journal of Raman Spectroscopy* 47(1):124–141
- Hoffman B (2003) The logic of suicide terrorism. *The Atlantic*. <http://www.theatlantic.com/magazine/archive/2003/06/the-logic-of-suicide-terrorism/302739/>, [Accessed 2 July 2016]
- Hoos H, Stützle T (2005) *Stochastic Local Search. Foundations and Applications*. Morgan Kaufmann
- Kaplan EH, Kress M (2005) Operational effectiveness of suicide-bomber-detector schemes: A best-case analysis. *Proceedings of the National Academy of Sciences of the United States of America* 102(29):10,399–10,404
- Kim M, Batta R, He Q (2016) Optimal routing of infiltration operations. *Journal of Transportation Security* 9(1):87–104

- Kroenig M, Pavel B (2012) How to deter terrorism. *The Washington Quarterly* 35(2):21–36
- Şenses Kurukız H (2016) Istanbul airport terror attack death toll rises to 44. *Anadolou Agency*. <http://aa.com.tr/en/todays-headlines/istanbul-airport-terror-attack-death-toll-rises-to-44/599383>, [Accessed 2 July 2016]
- Larrañaga P, Lozano JA (2002) Estimation of distribution algorithms: A new tool for evolutionary computation, vol 2. Springer Science & Business Media
- Mühlenbein H, PaaßG (1996) From recombination of genes to the estimation of distributions I. Binary parameters. In: Voigt H, Ebeling W, Rechenberg I, Schwefel H (eds) *Parallel Problem Solving from Nature - PPSN IV*, Springer-Verlag, Berlin Heidelberg, Lecture Notes in Computer Science, vol 1141, pp 178–187
- Neri F, Cotta C, Moscato P (eds) (2012) *Handbook of Memetic Algorithms*, Studies in Computational Intelligence, vol 379. Springer-Verlag, Berlin Heidelberg
- Nie X, Batta R, Drury CG, Lin L (2007) Optimal placement of suicide bomber detectors. *Military Operations Research* 12(2):65–78
- NRC (2004) Existing and potential standoff explosives detection techniques. National Research Council of the National Academies, Natl. Acad. Press, Washington, DC
- Rand Corporation (2009) RAND database of worldwide terrorism incidents. [Data File] http://smapp.rand.org/rwtid/search_form.php
- Shaffer JP (1986) Modified sequentially rejective multiple test procedures. *Journal of the American Statistical Association* 81(395):826–831
- Singh S (2007) Sensors—an effective approach for the detection of explosives. *Journal of Hazardous Materials* 144(1–2):15–28
- Yan X, Nie X (2016) Optimal placement of multiple types of detectors under a small vessel attack threat to port security. *Transportation Research Part E: Logistics and Transportation Review* 93:71–94
- Yinon J (2007) *Counterterrorist Detection Techniques of Explosives*. Elsevier Science B.V., Amsterdam