Mathematical optimization for social distancing

Martina Fischetti[(1)], Matteo Fischetti[(2)], Jakob Stoustrup[(3)]

[(1)] Vattenfall BA Wind, København, Denmark:
martina.fischetti@vattenfall.com

[(2)] DEI, University of Padova, Italy: matteo.fischetti@unipd.it

[(3)] Automation & Control, Department of Electronic Systems,
Aalborg University, Denmark: jakob@es.aau.dk

May 2020

Abstract

The spread of viruses such as SARS-CoV-2 brought new challenges to our society, including a stronger focus on safety across all businesses. In particular, many countries have imposed a minimum social distance between people in order to ensure their safety. This brings new challenges to many customer-related businesses, such as restaurants, offices, etc., on how to located their facilities under distancing constraints. We propose a parallelism between this problem and the one of locating wind turbines in an offshore area. Even if the two problems may seems very different, there are many analogies between them. In particular, both problems require fitting facilities (turbines or customers) in a given area while ensuring a minimum distance between them. Similarly to nearby customers who can infect each other, also nearby turbines "infect" each other by casting wind shadows (the so-called "wake effect") that cause production losses. In both problems we want to minimize the interference/infection. The solution of both problems will therefore favor solutions where facilities are as spread as possible. The discovery of this parallelism between the two applications allowed us to apply a recently published approach for wind-farm layout design to produce optimized facility layouts subject to social distancing constraints as those arising in the time of COVID-19 pandemic. These methods allow us to challenge the current (manual) layouts and provide new insights on how to improve them. In particular we show that optimized layout are far from trivial to design and that Mathematical Optimization can make an impact, helping businesses while ensuring safety.

Keywords: Social distancing, Mathematical Optimization, Mixed Integer Linear Programming, COVID-19

1 Introduction

In this paper we study the optimal positioning of people in a common area in order to minimize the spread of viruses such as SARS-CoV-2. The re-opening of many countries after lockdown, indeed, puts new challenges for many businesses as they want to ensure safety of their customers and personnel, while still making a profit and offering their usual services. Many countries impose a minimum distance between people in order to ensure safety, the so-called social distancing constraint. This means that shared spaces (like offices, restaurants, public transports, etc.) are now facing a new challenge that concerns fitting their usual customers while respecting social distancing restrictions.

The challenge can be interpreted in different ways: on one side, we may want to place as many "facilities" (e.g., customers) as possible in a given area while fulfilling the social distancing requirement. One example can be a restaurant where we want to fit as many tables as possible to ensure a profitable business and customer satisfaction, while not putting customers at risk. On the other side, we may have a fixed number of facilities that we want to fit in the most safe way in a given area. One example can be a restaurant with a limited personnel capacity and a large room available for tables, where the strategy is to fit a predefined number of tables while maximizing the safety of the clients.

The intuition would tell us that, placing facilities at the maximum distance one from each other would ensure the most safe layout. But is this enough? Also, the challenge of the optimal placement of facilities is often solved manually, by placing them at a fixed distance from one another. Is this a good strategy?

Our experience on a different, yet similar, problem (the wind farm layout optimization problem), allowed us to verify these intuitions through computational experiments, thus bringing new light on the structure of the optimal solutions under different constraints and objectives.

In particular, in this paper we will challenge the intuition that having a regularly-spaced layout is an optimal choice, looking both at the variant of the problem where we maximize the number of facilities, and at the version aimed at the minimization of virus spread. We will do so by explicitly modelling (in an admittedly simplified way) a possible spread function for the virus, and by letting a Mathematical Optimization algorithm decide the best placement of facilities. We will show that the optimized placement is often far from trivial, and depends on the geometrical shape of the available area. We will see from our results, for example, that positioning as many facilities as possible on the borders of the available area tends to be a winning strategy, while the positions in the center of the area should be avoid as much as possible as they are surrounded in all directions by many other (potentially infectious) facilities. Whereas this might seem unsurprising, we shall also demonstrate that the specific topologies can be non-trivial.

Our ultimate goal is to show how the usage of sound Mathematical Optimization algorithms can improve the layout of public spaces while decreasing the spread of viruses.

2 Optimal social distance and optimal wind farm layout

The professional expertise of the first author concerns the usage of optimization in the design phase of wind farms. Among other topics, she worked on the offshore wind farm layout optimization problem, which consists in deciding where to place turbines in a given area in order to maximize production while reducing costs [8, 4, 9, 6]. But what is the relationship between wind farms and social distancing? Let us introduce the reader to the wind farm layout optimization problem first, so that the similarities with the problem at hand will result more clearly.

The offshore wind energy business in Europe is based on an auction system where a country puts a well-defined offshore area on tender, and different companies compete to get the rights to construct the wind farm in that area. The company that can provide energy at the lower cost will get the right to construct and operate the new offshore wind farm. Therefore each company wants to design the new wind farm minimizing costs and increasing profitability. This task typically involves optimizing the placement of turbines within the given area. A key aspect in the optimization is to take wake effects into account. The wake effect is the interference phenomenon for which, if two turbines are located one close to another, the upwind turbine creates a shadow on the one behind, see e.g. [11]. This is of great importance in the design of the layout since it results in a loss of power production for the turbines downstream, that are also subject to a strong (hence damaging) turbulence. Also, nearby wind farms might need to be considered in the optimization as they can also interfere with the new one.

Figure 1: Wake effect on a real wind farm (Horns Rev 1): regular layouts like the one in the picture can be very inefficient for certain specific wind scenarios, resulting into a greatly reduced energy production for the whole wind farm. [Source: Vattenfall]

Given a certain wind scenario, defined by the wind intensity and direction, the wake effect can be simply modelled as a cone, centered in the upwind turbine, that fades away with distance (Jensen's model [10]). Nevertheless, when the wind farm layout problem is solved during the design phase, one needs to consider the full variability of the wind. As a matter of fact, considering only one wind scenario would create a layout which is optimal for that specific wind condition, but potentially highly suboptimal when the wind changes intensity and/or direction. In [8] a method to consider the full variability of the wind when optimizing the layout, is proposed. Intuitively, this method lets the optimizer consider a weighted intersection of the interference cones for all possible wind scenarios, where more likely scenarios have a higher weight in the combination. Visually, this means that the interference becomes a star of cones around the turbine. In practical applications with thousands of possible wind scenarios, this star is so dense that it looks like a continuous interference shade around the turbine, with more intense values in the main wind directions; see Figure 2 for an illustration.

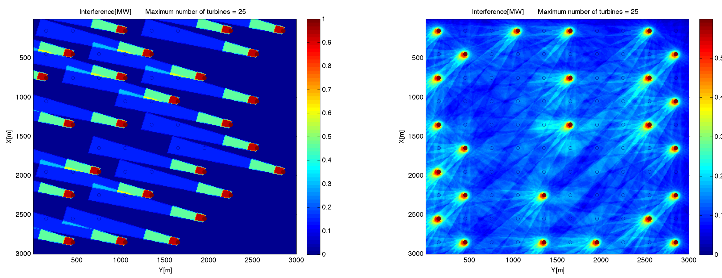

Figure 2: The interference between turbines can be visualized as a cone when only one wind scenario is considered (left plot). In the design of an offshore wind farm, however, the full variability of the wind needs to be considered, as visualized in the right-hand-side plot. [Source: [3]]

In practical applications, a minimum and/or maximum number of turbines to be located in the area can be imposed, together with a minimum distance between turbines (to avoid the blades clash, and also for turbulence considerations). There can be obstacles within the offshore area (such as natural reserves, preexisting infrastructures, bad seabed areas etc.), which are areas where the turbines cannot be located. The wind farm layout optimization problem therefore consists in locating a given number of turbines (or as many as profitable) in a given area, ensuring a minimum distance between turbines and minimizing the interference between them.

What about our problem about social distancing in a public place such as a restaurant? It also consists in locating a given number of facilities (tables or customers) in a given area, ensuring a minimum distance between them (legal or recommended social distance) and minimizing the potential virus spread between facilities.

In the wind farm layout problem the interference is a well-studied function that represents the loss of power due to wake effect. In the social distancing problem, instead, the interference is a function that represents the spread of the virus. Both interference interpretations depend on the distance—the further the better. As we will see in more details in Section 4, we can already intuitively see the similarity between the interference function between wind turbines (second plot of Figure 2) and an infection function around customers (as, for example, the one in the top-right plot of Figure 5) .

Our experience on the wind farm layout problem shows that optimal layouts tend to use the borders of the available area, where the turbines create less interference to other turbines. It also shows that traditional manual layouts where turbines were placed on a regular grid, are highly sub-optimal as significantly higher production can be achieved by a smarter (i.e., optimized) placement of turbines. The resolution of the problem in practical applications is far from trivial, and state-of-the-art Mathematical Optimization techniques have proved to make a huge impact in the practical resolution of the problem [7].

Can we then use our wind farm expertise to provide private and public businesses with insights on how to optimally place customers, in order to minimize the spread of the virus? Is it true also for this problem, that regular layouts are often suboptimal?

3 Optimization model (for wind farms)

The first author proposed in [8] a Mixed Integer Linear Programming (MILP) model for the wind farm layout problem. The overall area available is sampled to define a discrete set V of possible positions, and a binary variable x_i is defined for each $i \in V$, taking value 1 if and only if a turbine is built at position $i \in V$.

The optimizer considers:

a) a minimum and maximum number of turbines that can be built;

b) a minimum separation distance between any pair of turbines, to ensure that the blades do not physically clash;

c) the interference between installed turbines (wake effect).

Let

- I_{ij} be the interference (loss of power) experienced by position j when a turbine is installed at position i, with $I_{jj} = 0$ for all $j \in V$; Jensen's model [10] can be used to compute such an interference;
- P_i be the power that a turbine would produce if built (alone) at position i;
- $N_{\min}$ and $N_{\max}$ be the minimum and maximum number of turbines that can be built, respectively;
- d_{ij} be the distance between positions i and j;
- $D_{\min}$ be the minimum distance required between two turbines.

In addition, let $G_I = (V, E_I)$ denote an "incompatibility" undirected graph with

$$E_I = \{ [i,j] : i, j \in V,\, d_{ij} < D_{\min},\, i < j \}.$$

A natural quadratic objective function (to be maximized) for our problem reads

$$\sum_{i \in V} (P_i x_i - (\sum_{j \in V} I_{ij} x_j)\, x_i) \tag{1}$$

and can be restated as

$$\sum_{i \in V} (P_i x_i - w_i) \tag{2}$$

where the continuous variable w_i is defined as

$$w_i = \Big(\sum_{j \in V} I_{ij} x_j \Big) x_i = \begin{cases} \sum_{j \in V} I_{ij} x_j & \text{if } x_i = 1; \\ 0 & \text{if } x_i = 0 \end{cases}$$

and denotes the total interference caused by a turbine built in position i. The MILP model then reads

$$\begin{aligned}
\max \quad & z = \textstyle\sum_{i \in V} (P_i x_i - w_i) && && (3)\\
\text{s.t.} \quad & N_{\min} \le \textstyle\sum_{i \in V} x_i \le N_{\max} && && (4)\\
& x_i + x_j \le 1, && [i,j] \in E_I && (5)\\
& \textstyle\sum_{j \in V} I_{ij} x_j \le w_i + M_i(1 - x_i), && i \in V && (6)\\
& x_i \in \{0, 1\}, && i \in V && (7)\\
& w_i \ge 0, && i \in V. && (8)
\end{aligned}$$

The objective function (3) maximizes the total power production by taking interference losses into account. Constraints (4) impose a minimum and a maximum number of turbines that can be constructed in the area. If $N_{min} = N_{max}$, then we are actually imposing to build a fixed number of turbines, otherwise the optimizer can define the best number of turbines (within N_{min} and N_{max}) to be located. Constraints (5) ensure the minimum distance between turbines. Constraints (6) force the correct value for variables w_i; here, a big-M term $M_i \gg 0$ is used to deactivate the constraint in case $x_i = 0$, namely

$$M_i = \sum_{j \in V : (i,j) \notin E_I} I_{ij}.$$

Finally (7) and (8) define our binary and continuous variables, respectively. As shown in details in [8], using a single index variable w_i allows this model to solve larger instances compared with equivalent two-index models in the literature (e.g., [1, 2]). Another strength of this formulation is the ability of easily dealing with multiple wind scenarios; the reader is again referred to [8] for further details.

Based on the above model, a sound matheuristic [5] solution approach has been developed in [8].

4 Modelling virus spreading

Wake effects between turbines have been largely studied in the wind energy literature, and well-established interference models exist in the literature [10]. As far as we know, however, the same does not hold for virus-spread functions, in the sense that there is not widely-accepted and sufficiently simple models. For recent CFD based models that takes ventilation air-flows into account please refer to [12] and references there-in. The complexity, however, of CFD models prohibits a feasible optimization setup.

Therefore we decided to analyse alternative functions that can be used to define the interference matrix (I_{ij}) used in our optimization model. These functions are just illustrative and should be considered with caution as they are not based on a careful pathogenic spread analysis or validated e.g. through tracer gas experiments. Our aim is just to find a function that favors layouts where the facilities are scattered among the available area even with $D_{\min} = 0$, i.e., even without imposing explicitly a minimum distance between facilities.

When I_{ij} has been defined, we assume that the "total virus exposure" (i.e., the infection risk) experienced by a customer sitting in position j is computed as

$$\sum_{i \in V, i \neq j} I_{ij} x_i,$$

where $x_i = 1$ if a source of virus (e.g., a potentially-infected customer) is present at position i, and $x_i = 0$ otherwise.

Let d_{ij} represent the Euclidean distance (in meters) between positions $i, j \in V$, and let $d_{\max}$ be the maximum such distance. We consider the following alternative definitions for the interference I_{ij} that a facility located at position i causes to a customer in position j at a distance $d_{ij} \geq 10^{-4}$; by convention, $I_{ij} = 0$ whenever $d_{ij} < 10^{-4}$, which implies $I_{ii} = 0$ for all $i \in V$:

$$I_{ij} = d_{max} - d_{ij} \tag{9}$$

$$I_{ij} = e^{-d_{ij}^2/2} \tag{10}$$

$$I_{ij} = 1/d_{ij} \tag{11}$$

$$I_{ij} = 1/d_{ij}^2 \tag{12}$$

$$I_{ij} = 1/d_{ij}^3 \tag{13}$$

Definition (9) considers an infection risk inversely proportional to the distance. At first glance, this would seem a reasonable assumption, but it turns out to be a very bad choice: even when positioning just 3 facilities on a line segment, an optimal solution will position two facilities and the endpoints of the segment, while the third one can freely be located at any other point of the segment without affecting the solution value—while we would expect the optimal solution be unique and place the third facility on the middle point of the segment.

Definition (10) assumes the infection risk can be modeled as a Gaussian function with variance $\sigma^2 = 1$. This can be a realistic assumption to model the trajectory of droplets, which are expected to decay rapidly and fall down within 1-2m or so.

Definitions (11) to (12) assume instead that a virus aerosol is spread uniformly on a 1- or 2- or 3-dimensional space, respectively. Indeed, in a k-dimensional space a sphere of ray d_{ij} centered at position i (where the virus aerosol is emitted) has a volume proportional to d_{ij}^k. Thus, assuming a uniform virus distribution within the sphere, the virus concentration at position j is proportional to $1/d_{ij}^k$. In this view, definition (11) seems a realistic model only for 1-dimensional problems where virus is spread along a line, while (12) and (13) seems more suited for the 2- and 3-dimensional cases (the latter arguably being the most realistic one).

4.1 Computational study of alternative virus-interference functions

As already discussed, we are interested in defining the interference matrix (I_{ij}) in a way that favors the spread of facilities on the available area. In the next subsections, we will therefore consider regular 1- and 2-dimensional areas and compare the optimised layouts resulting from the alternative interference definitions. In those experiments, we fixed the number k of facilities to be located (by setting $N_{\min} = N_{max} = k$), we defined $P_i = 0$ for all $i \in V$ (no power production), and we did not require any minimum distance between facilities ($D_{\min} = 0$). In this way, we only evaluated the effect of the alternative interference-matrix definitions in producing scattered layouts.

4.1.1 Positioning facilities on a line

In Figure 3 we consider the problem of positioning 10 facilities on a line segment, using the alternative interference definitions (9) to (13). As expected, the two endpoints of the segment are selected in all cases—using the border of the available area turns out to be a successful policy in wind farm design as well. The layout in the top subfigure (definition (9)) is completely unsatisfactory in terms of social distances, as the selected points define two clusters of five points each, located beside the line endpoints. The second layout based on the Gaussian function (10) is more regular, but still uses two pairs of almost-overlapping positions beside the line endpoints. The remaining layouts are obtained with (11) to (13), and allow for a satisfactory distancing among facilities.

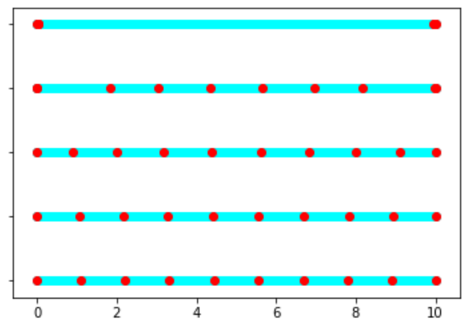

Figure 3: Positioning 10 facilities on a line using interference definitions (9) to (13) (from top to bottom) .

4.1.2 Positioning facilities within a square

In Figure 4 we consider the problem of positioning 20 facilities within a square. As expected, the border points are very attractive and the four edges contain most of the built facilities. As in the 1-dimensional case, definition (9) produces clusters of facilities on the four vertices of the square (top-left subfigure). Using the Gaussian function (10) produces the more satisfactory layout of the top-right subfigure, while the bottom-left subfigure refers to (11). The most satisfactory layouts are obtained using definition (12) (bottom-right subfigure) and (13) (not reported as identical to the one in the bottom-right subfigure).

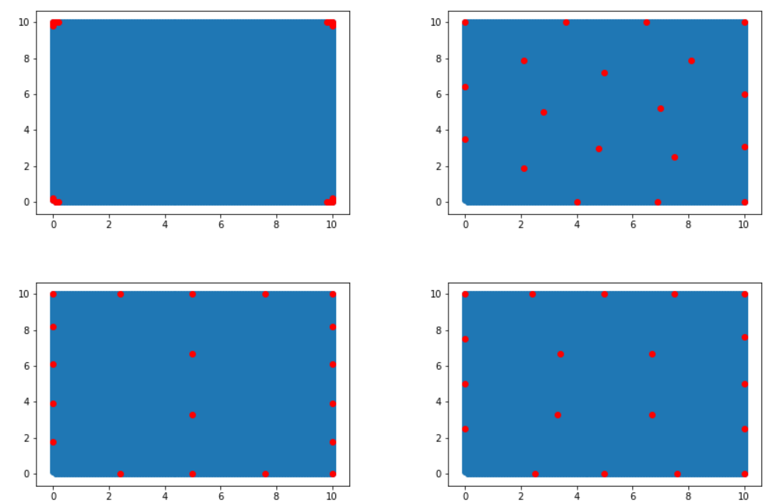

Figure 4: Positioning 20 facilities within a square area using definitions (9) (top-left) to (12) (bottom-right); the layout using definition (13) is the same as the one reported in the bottom-right subfigure.

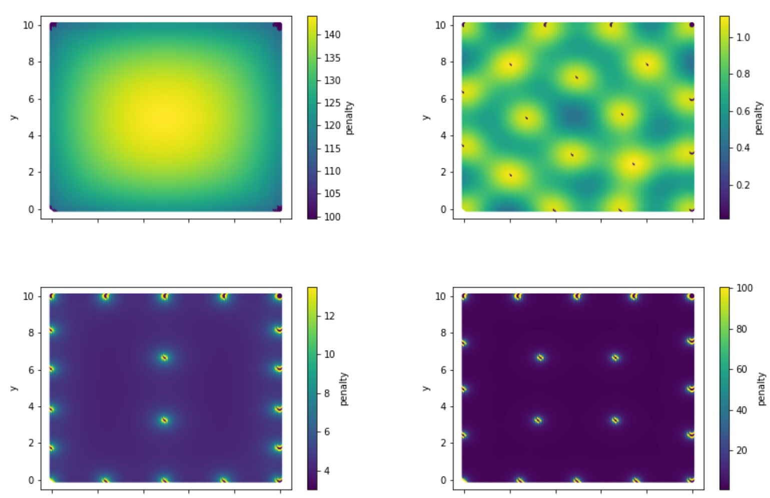

Figure 5: Penalty measuring the infection level (i.e., the total interference) using the interference definitions (9) (top-left subfigure) to (12) (bottom-right subfigure). Penalty scales are different in each subfigure. The plot using definition (13) is very similar to the one reported in the bottom-right subfigure, hence it is omitted.

Figure 5 plots the total interference $\sum_{i \in V, i \neq j} I_{ij} x_i$ perceived by each point $j \in V$, by using alternative interference definitions. In the virus-spread context, this is a measure of the probability of infection that depends on the location of all the built facilities. The top-left subfigure confirms once again that definition (9) is not adequate to represent virus spread, as the probability of infection is almost uniform in the square—the total interference ranges from about 140 in the center, and about 100 in the four corners where the facilities are actually clustered.

On the whole, it seems that definition (13) qualifies as the most reliable modelling formula of virus-aerosol spreading in an open-space environment—among the considered ones—hence we will use it in the forthcoming experiments.

5 Applications

We next address some practical examples of optimized facility layout under social distancing constraints.

5.1 Restaurants

Let us suppose we own a restaurant which has a certain space to place customer tables (an indoor room, or a given outdoor space). This space could have any shape and potentially have areas where tables cannot be placed (access to security exits, structural elements in the room, outdoor trees or fountains, etc.). We have to ensure a minimum distance between tables, as imposed by the local government regulations, and at the same time we want to fit as many tables as possible—even one more table can make a difference in the final income for the day. Also, among the feasible layouts we would like to choose one that minimizes the infection probability of the customers, measured though a suitable "virus spread" function—we used definition (13) in our experiments. Most restaurants are today solving this optimization challenge manually, often by placing tables aligned in rows and at a regular distance one from one another. As we will see, optimized solutions for this problem are not as regular as the manual ones—positioning tables in a less-regular but smarter way can satisfy more customers and reduce virus spread.

We considered two real cases to test the impact of optimizing table location in the restoration business. Our two examples are a cafeteria in Italy (Bar Nazionale, Padua) and a brewpub in Denmark (Brus, Copenhagen). Both businesses have an available outdoor area where they can place tables. Using online satellite views (first plot of Figures 6 and 7), we identified the available area for tables (red area in the second plot of Figures 6 and 7). Note, for example, that in the case of Brus (Figure 7) the presence of a tree does not allow to place tables in the middle part of the area, which was therefore excluded from the set of available positions. Being able to easily exclude areas from the optimization also allows for the definition of free corridors for customer movements, as in Figure 6.

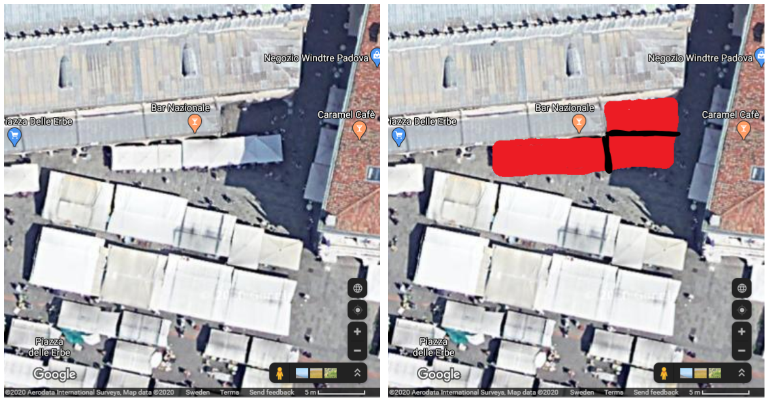

Figure 6: Example of optimization of table placement for the outside area of a cafeteria in Northern Italy. The available area for placing tables is highlighted in red in the second plot.

Figure 7: Example of optimization of table placement for the outside serving area of a brewpub in Denmark. The available area for placing tables is highlighted in red in the second plot.

5.1.1 Fitting more tables under social distancing rules

The first problem we would like to solve is to place as many tables as possible in the available area, while of course complying with the minimum distance between tables imposed by country's regulations. In our model, this is easily obtained by setting $N_{\min} = 0$, $N_{\max} = +\infty$, $D_{\min}$ to the required value (e.g., 3m center-to-center), and all P_i's to a large positive value (10^3 in our experiments) to favor building as many facilities as possible.

A naïve way to manually solve this problem is to locate tables on a regular grid, starting from one angle of the available area and locating each new table at the given minimum distance. Let us use this approach for the Italian Bar Nazionale of Figure 6: we define a regular grid starting from the right bottom corner and we put tables regularly at 3m distance. The result (and the underlying 3x3 grid) is visualized in the first plot of Figure 8. In this way the restaurant can locate 7 tables.

Then we try a more sophisticated approach, namely, we solve the table layout problem using our optimization method. The result is shown in the second plot of Figure 8: the optimizer could locate 10 tables, while still satisfying the minimum distance of 3m.

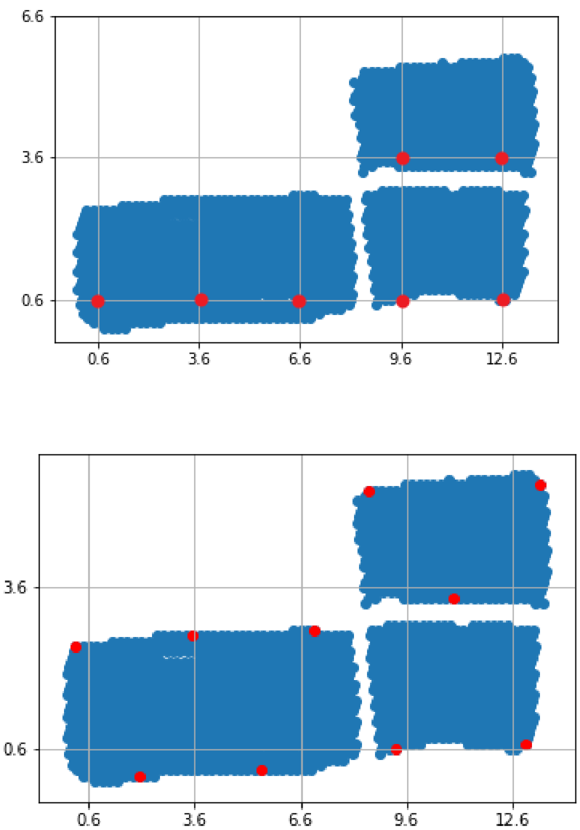

Figure 8: Example of optimization of table placement for the outside area of cafeteria “Bar Nazionale” in Northern Italy. The tables are placed following a manual approach (based on a super-imposed 3m × 3m regular grid) in the first plot. The second plot shows the optimized placement of tables, still with minimum distance of of 3m, using an optimization algorithm: 3 more tables can be fit in the same area.

We will now address our second test case, namely, the Brus brewpub in Copenhagen of Figure 7. The manual layout is defined here by imposing a regular grid starting from the top corner of the available area (shown on the plots of Figure 9). The manual approach could locate 30 tables (red dots in the first plot of Figure 9), while our optimization method could locate 36 turbines (red dots in the second plot of Figure 9). Having 6 more tables increases the capacity for customers of 20%: this can make a significant difference in terms of daily profits of the brewpub, without impacting the compliance to the local social distancing rules (here assumed to be 3m) and the safety of the customers.

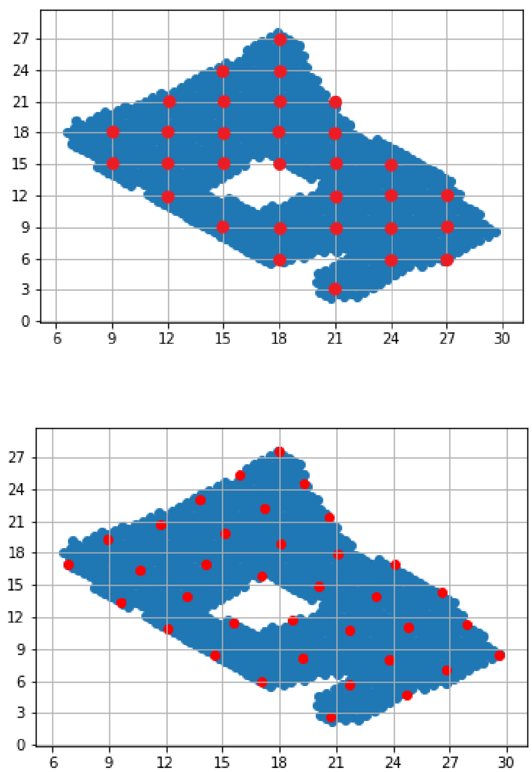

Figure 9: Example of optimization of table placement for the outside area of "Brus" brewpub in Denmark. Tables are placed following a manual approach (based on a super-imposed 3m × 3m regular grid) in the first plot. The second plot shows the optimized placement of tables at a minimum of 3m distance, using an optimization tool: 6 more tables can be located.

It is clear from these tests that the optimal placement of tables is often not straightforward, and that the usage of optimization methods can significantly increase the capacity of a restaurant to fit customers, and thus its daily revenue.

5.1.2 Fitting a given number of tables while minimizing virus spread

Another variant of the problem consists in fitting a fixed amount of tables in the area, while maximizing the safety of the customers. For these tests, we will fit as many tables as the manual solutions would fit (see first plot of Figures 8 and 9) but in a safer way. In other words, with respect to the previous tests, the focus is now shifted from maximizing the table number to minimizing the virus spread, while still fulfilling the country regulations on minimum table distances.

With respect to our two test cases we will therefore still require the minimum distance of 3m between tables, but we imagine that we want to locate 7 tables for the Italian cafeteria (though we have seen that up to 10 tables could fit in the same area) and 30 tables for the Danish brewpub (while have seen that up to 36 tables could fit as well). This is simply obtained in our model by setting $N_{\text{min}} = N_{\text{max}}$ (= 7 or 30). We use the virus-spread functions (13) to measure the risk of infection between tables. The resulting layouts are shown in Figure 10.

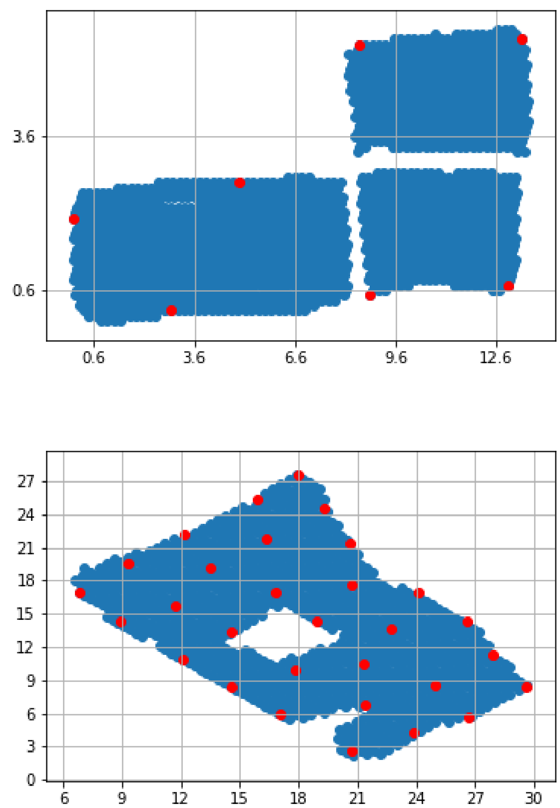

Figure 10: Minimizing virus spread for a fixed number of facilities smaller than the maximum capacity.

5.2 Beach umbrellas

Another possible application of our optimization method is the optimal location of umbrellas on a beach. Many seaside activities are indeed challenged by the COVID-19 restrictions, and owners of seaside areas can face difficulties in secure their income while ensuring safety. In countries like Italy, many beach areas are managed by privates, who rent beach facilities (such as umbrellas, sunbeds, chairs, etc) to customers. Due to COVID-19, social distance limitations also apply in defining the exact position of the umbrellas on the beach. Using optimization methods instead of relaying on manual layout, can have a big impact also in this case, allowing one to fit more customers while not compromising on their safety. For example, we considered a real case from beach "Bagni Alberoni", located in Venice, Italy. The first plot of Figure 11 shows the actual layout designed by the beach owners to cope with the required 4m minimum distance between umbrellas. This solution locates 203 umbrellas in the available area. We gave the same area (blue in the second plot of Figure) on input to our optimizer, together with the minimum distance of 4m. The interference definition (13) was considered. Our goal was to fit as many umbrellas as possible within the given limitations. Our optimizer was able to fit 211 umbrellas–having 8 more umbrellas to rent out over the summer season, can make a big economical impact for local business.

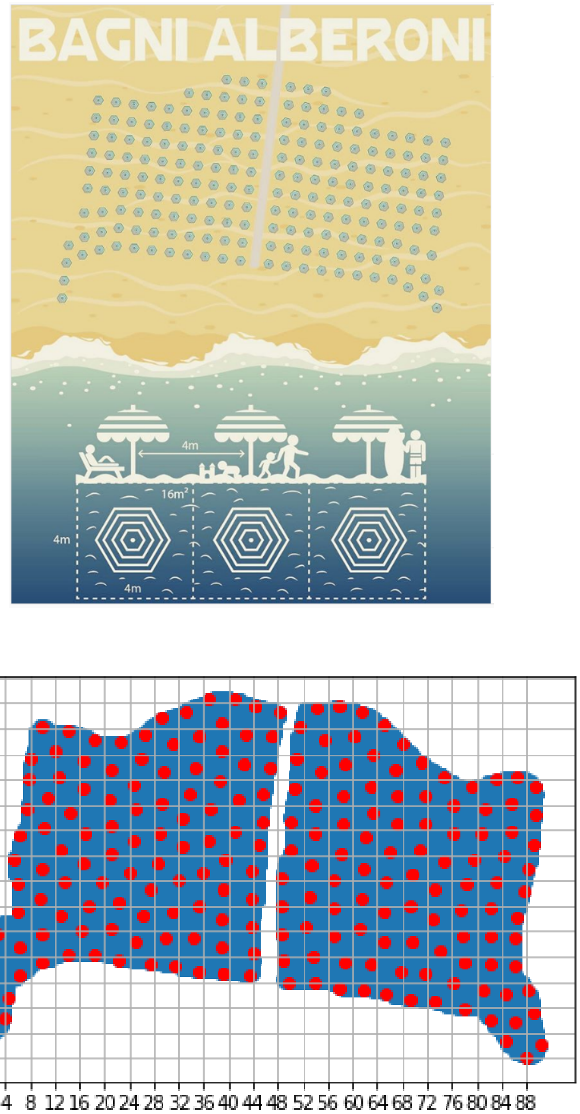

Figure 11: Example of a real case—the Venice beach "Bagni Alberoni". Beach umbrellas must ensure a minimum distance of 4m (center-to-center). The manual solution actually implemented (top) allocates 203 beach umbrellas, while the optimized one (bottom) is able to fit 211 beach umbrellas using a less-regular pattern.

6 Conclusions and future work

In this paper we have studied the problem of locating facilities in a given area, subject to social distancing constraints as those arising in the time of COVID-19. We have proposed a parallelism between this problem and the one of locating wind turbines in an offshore area, which allowed us to apply state-of-the-art solution approaches for the latter problem to produce optimized facility layouts. We have analyzed alternative definitions of the interference function used to model virus spread, and have compared them on simple cases. Then we have addressed possible applications of our optimization technology, showing that improved solutions can be obtained with less-regular (but more efficient) layout patterns than those typically found manually.

Future work can address different applications such as seat allocation in transport systems (buses, trains or airplanes).

Another challenging research direction is the definition of an interference function that models virus spread along specific directions (rather than uniformly over the surrounding space), e.g., because of the presence of air conditioning and/or Plexiglas settings. This would produce very irregular interference matrices (akin to those arising in offshore wind farms with predominant wind scenarios) that are very difficult to handle by a manual solution approach, thus making our methodology even more appealing.

References

[1] Archer, R., Nates, G., Donovan, S., Waterer, H.: Wind turbine interference in a wind farm layout optimization – mixed integer linear programming model. Wind Engineering **35, 2**, 165–178 (2011)

[2] Fagerfjall, P.: Optimizing wind farm layout - more bang for the buck using mixed integer linear programming. Master's thesis, Department of Mathematical Sciences, Chalmers University of Technology and Gothenburg University, Göteborg, Sweden (2010)

[3] Fischetti, M.: Mathematical programming models and algorithms for offshore wind park design. Ph.D. thesis, Technical University of Denmark, DTU Management (2018)

[4] Fischetti, M.: Improving profitability of wind farms with operational research. Impact **2019**(1), 30–34 (2019)

[5] Fischetti, M., Fischetti, M.: Matheuristics. In: R. Martí, P.M. Pardalos, M.G.C. Resende (eds.) Handbook of Heuristics, pp. 121–153. Springer International Publishing, Cham (2018)

[6] Fischetti, M., Fraccaro, M.: Machine learning meets mathematical optimization to predict the optimal production of offshore wind parks. Computers & Operations Research **106**, 289–297 (2019)

[7] Fischetti, M., Kristoffersen, J.R., Hjort, T., Monaci, M., Pisinger, D.: Vattenfall optimizes offshore wind farm design. INFORMS Journal On Applied Analytics **50**(1), 80–94 (2020)

[8] Fischetti, M., Monaci, M.: Proximity search heuristics for wind farm optimal layout. Journal of Heuristics **22**(4), 459–474 (2016)

[9] Fischetti, M., Pisinger, D.: Mathematical optimization and algorithms for offshore wind farm design: An overview. Business & Information Systems Engineering **61**(4), 469–485 (2019)

[10] Jensen, N.: A note on wind generator interaction. Tech. rep., Riso-M-2411(EN), Riso National Laboratory, Roskilde, Denmark (1983)

[11] Odgaard, P.F., Stoustrup, J.: Fault tolerant wind farm control - a benchmark model. In: Proceedings of the 2013 IEEE Multi-conference on Systems and Control, pp. 412–417. Hyderabad, India (2013). DOI 10.1109/CCA.2013.6662784

[12] Villafruela, J., Olmedo, I., San José, J.: Influence of human breathing modes on airborne cross infection risk. Building and Environment **106**, 340 – 351 (2016). DOI 10.1016/j.buildenv.2016.07.005