Comparing Parallel Surrogate-based and Surrogate-free Multi-Objective Optimization of COVID-19 vaccines allocation

Guillaume Briffoteaux, Romain Ragonnet, Pierre Tomenko, Mohand Mezmaz, Nouredine Melab, Daniel Tuyttens

Comparing Parallel Surrogate-based and Surrogate-free Multi-Objective Optimization of COVID-19 vaccines allocation

Guillaume Briffoteaux[1,3], Romain Ragonnet[2], Pierre Tomenko[1], Mohand Mezmaz[1], Nouredine Melab[3] and Daniel Tuyttens[1]

[1] Mathematics and Operational Research Department, University of Mons, Belgium
{guillaume.briffoteaux,pierre.tomenko,mohand.mezmaz,daniel.tuyttens}@umons.ac.be
[2] School of Public Health and Preventive Medicine, Monash University, Australia
romain.ragonnet@monash.edu
[3] University of Lille, Inria, UMR 9189 - CRIStAL, France
nouredine.melab@univ-lille.fr

Abstract. The simulation-based and computationally expensive problem tackled in this paper addresses COVID-19 vaccines allocation in Malaysia. The multi-objective formulation considers simultaneously the total number of deaths, peak hospital occupancy and relaxation of mobility restrictions. Evolutionary algorithms have proven their capability to handle multi-to-many objectives but require a high number of computationally expensive simulations. The available techniques to raise the challenge rely on the joint use of surrogate-assisted optimization and parallel computing to deal with computational expensiveness. On the one hand, the simulation software is imitated by a cheap-to-evaluate surrogate model. On the other hand, multiple candidates are simultaneously assessed *via* multiple processing cores. In this study, we compare the performance of recently proposed surrogate-free and surrogate-based parallel multi-objective algorithms through the application to the COVID-19 vaccine distribution problem.

1 Introduction

In this paper, we address a multi-objective (MO) COVID-19 vaccines allocation problem. We aim to identify vaccines allocation strategies that minimize the total number of deaths and peak hospital occupancy, while maximizing the extent to which mobility restrictions can be relaxed. The onset of the COVID-19 outbreak has been rapidly followed by the development of dedicated simulation software to predict the trajectory of the disease [1, 2]. The availability of such tools enables one to inform authorities by formulating and solving optimization problems. In [3], a SEIR-model (Susceptible, Exposed, Infectious, Recovered) is deployed to simulate COVID-19 impacts. A single-objective (SO) problem is subsequently derived and handled by grid-search to regulate the alleviation of social restrictions. Multiple SO optimizations are carried out independently by

a simplex or a line search algorithm in [4–6] to efficiently allocate doses of vaccines to the age-categories of a population. The number of infections, deaths and hospital admissions are considered as the possible objective. The prioritization rules approved by the government of the studied cohort are integrated as constraints in the linear programming model presented in [7] to minimize mortality. In [8], multiple indicators are combined into a scalar-valued objective function. The MO formulation exhibited in [9] consists in maximizing the geographical diversity and social fairness of the distribution plan. Nevertheless, the MO problem is scalarized into a SO one that is then solved by a simplex algorithm. The approach used by Bubar and colleagues [10] is significantly different to ours, as the authors predefined a set of vaccination strategies and selected the most promising approach among them. In contrast, our continuous optimisation approach automatically designs strategies in a fully flexible way. The optimisation problem solved by McBryde and colleagues [11] is closer to that presented in our work since a similar level of flexibility was allowed to design optimal vaccines allocation plans. However, the authors used a simpler COVID-19 model resulting in significantly shorter simulation times, such that optimisation could be performed using more classical techniques. To the best of our knowledge, it has not been suggested yet to simultaneously minimize the number of deaths, peak hospital occupancy and the degree of mobility restriction through a MO-formulated problem. The fact that we consider the level of restrictions as one of the objectives to minimise represents a novelty compared to the previous works.

Despite the relative computational expensiveness of infectious disease transmission simulators, surrogate-based optimization has been rarely applied to the field. In [12], we harnessed surrogate models to determine the allocation of preventive treatments that minimize the number of deaths caused by tuberculosis in the Philippines. The identification of the regime for tuberculosis antibiotic treatments with lowest time and doses is formulated as a SO problem in [13] and solved by a method relying on a Radial Basis Functions surrogate model. The work presented in [14] deviates from this present study in that it aims to conceive a model prescribing the actions to perform according to a given situation. It is actually more related to artificial neural network hyper-parameters and architecture search. What is called "surrogate" in [14] is actually denominated "simulator" in simulation-based optimization. In this work, we combine machine learning and parallel computing to solve the MO vaccine distribution problem.

This study demonstrates the suitability of parallel surrogate-based multi-objective optimization algorithms on the real-world problem of COVID-19 vaccines allocation. The COVID-19-related problem is detailed in Section 2 and the MO algorithms are exposed in Section 3. Both surrogate-based and surrogate-free parallel MO approaches are applied to the real-world challenge in Section 4 and empirical comparisons are realized. Finally, conclusions are drawn in Section 5 and suggestions for future investigations are outlined.

2 COVID-19 vaccine distribution problem

The vast vaccination programs implemented over the last year or so all around the world achieved dramatic reductions of COVID-19 hospitalizations and deaths [15]. However, access to vaccination remains challenging, especially for low- to middle-income countries that are not able to offer vaccination to all their citizens [16]. The problem we are concerned with consists in optimizing the age-specific vaccines allocation plan to limit the impact of the disease in Malaysia under a capped number of doses. The population is divided into 8 age-categories of 10-years band from 0-9 years old to 70+ years old and the impact is expressed in terms of total number of deaths and peak hospital occupancy.

The simulation is realized in three phases by the AuTuMN software publicly available in `https://github.com/monash-emu/AuTuMN/`. The simulator is calibrated during the first phase with data accumulated from the beginning of the epidemic to the 1st of April 2021. The second phase starts at this latter date and lasts three months during which a daily limited number of doses is shared out among the population. Relaxation of mobility restrictions marks the kickoff of the third phase in the course of which a new distribution plan is applied involving the same number of daily available doses as in phase 2.

Decision variables $x_i \in [0, 1]$ for $1 \leqslant i \leqslant 8$ and for $9 \leqslant i \leqslant 16$ represent the proportions of the available doses allocated to the 8 age-categories for phase 2 and phase 3 respectively. Variable $x_{17} \in [0, 1]$ expresses the degree of relaxation of mobility restrictions where $x_{17} = 0$ leaves the restrictions unchanged and $x_{17} = 1$ means a return back to the pre-covid *era*. The following convex constraints convey the limitation of the number of doses during phases 2 and 3:

$$\sum_{i=1}^{8} x_i \leqslant 1 \text{ and } \sum_{i=9}^{16} x_i \leqslant 1 \tag{1}$$

The three-objective optimization problem consists in finding $\boldsymbol{x}^*$ such that

$$\boldsymbol{x}^* = \underset{\boldsymbol{x} \in [0,1]^{17} \text{ s.t. } (1)}{\arg\min} (g_1(\boldsymbol{x}), g_2(\boldsymbol{x}), 1 - \boldsymbol{x}_{17}) \tag{2}$$

where $g_1(\boldsymbol{x})$ is the simulated total number of deaths and $g_2(\boldsymbol{x})$ the simulated maximum number of occupied hospital beds during the period.

3 Parallel Multi-Objective Evolutionary Algorithms

3.1 Variation Operators of Evolutionary Algorithms

Evolutionary Algorithms (EAs) are harnessed to deal with the COVID-19-related problem exhibited previously. In EAs, a population of solutions is evolved through cycles of parents selection, reproduction, children evaluation and replacement. EAs are chosen because they have proven their effectiveness on numerous multi-objective real-world problems [17] where the objective functions are black-box

as it is the case in our scenario. The constraint being convex and analytically verifiable, it is thus possible to design specific reproduction operators that directly generate feasible candidates. Assuming that every feasible solution can be reached, this technique has shown to be a reliable one [18].

The specific cross-over operator, called *distrib-X*, considers the two phases and the degree of relaxation independently. For two parents $\boldsymbol{x}$ and $\boldsymbol{y}$, the last decision variable for the two children $\boldsymbol{z}$ and $\boldsymbol{t}$ is set such that $z_{17} = x_{17}$ and $t_{17} = y_{17}$. Regarding the second phase, let I and J be a random partition of $\{1, \dots, 8\}$. For the age categories in I, $\boldsymbol{z}$ receives the proportion of vaccines from $\boldsymbol{x}$ ($z_i = x_i$ for $i \in I$). The remaining proportion of available doses at this step is $r = 1 - \sum_{i \in I} x_i$. For the age categories in J, the remaining proportion of doses is shared out according to the proportion allocated to the corresponding age categories in $\boldsymbol{y}$. In other terms, for $j \in J$, $z_j = \frac{r \cdot y_j}{\sum_{j \in J} y_j}$. A similar treatment is applied to the variables associated to the third phase. The second child $\boldsymbol{t}$ is generated with an analogous procedure, where the roles of the parents are reversed.

The specific mutation operator, denoted *distrib-M*, disturbs a decision variable randomly chosen with uniform probability for $\{1, \dots, 8\}$, $\{9, \dots, 16\}$ and $\{17\}$. The last decision variable is mutated by polynomial mutation [17]. For the remaining ones, two age categories of the same phase are randomly selected and a random amount of doses are transferred from the first category to the second one. Both *distrib-X* and *distrib-M* are inspired by [12].

The *intermediate* and the *2-points* cross-over operators [17] are also considered for the sake of comparison. The *intermediate* strategy combines parents by random weighting average, while the *2-points* operator distributes portions of parents to the children. The portions are defined by two points with the first one separating phase 2 and phase 3 and the second one located between phase 3 and the relaxation decision variable x_{17}.

3.2 Parallel Multi-objective Evolutionary Algorithms

The major challenge in multi-objective optimization is to balance convergence and diversity in the objective space. Convergence is related to the closeness to the Pareto Front (PF) [17]. The PF is the set of the overall best solutions represented in the objective space and the Non-Dominated Fronts (NDFs) are approximations of the PF. Diversity is indicated by an extended coverage of the objective space by the NDFs. Hereafter, we present four algorithms to set this trade-off.

The first algorithm considered in the comparison is the Non-dominated Sorting Genetic Algorithm (NSGA-II) [19]. Firstly, to promote convergence, solutions pertaining to better NDFs are better ranked. Secondly, to favor diversity, solutions composing the same NDF are distinguished by setting the promise as high as the crowding distance is high. The proposed sorting is employed at the selection and the replacement steps of the EA.

The second algorithm reproduced for the experiments is the Reference Vector guided Evolutionary Algorithm (RVEA) proposed in [20] to handle many-

objective problems. A set of reference vectors is introduced in order to decompose the objective space and to enhance diversity. New candidates are attached to their closest reference vector, thus forming sub-populations among which only one candidate is kept at the replacement step. The new angle penalized distance chooses adaptively the candidate to be conserved by favoring convergence at the beginning of the search and diversity at latter stages. It is worth noting that the population size may change during the search in RVEA due to the possibility of empty sub-populations. In cases of degenerated or disconnected PF a high number of sub-populations become empty and the NDF obtained at the end of the search may not be dense enough. In the RVEA* variant, an additional reference vectors set is used to replace the reference vectors that would correspond to empty sub-populations.

In surrogate-based optimization, the additional trade-off between exploitation and exploration is to be specified. Minimizing the predicted objective vectors (POVs) produced by the surrogate boosts exploitation of known promising regions of the search space. Conversely, maximizing the predictive uncertainty enhances exploration of unknown regions.

The third algorithm is the surrogate-based Adaptive Bayesian Multi-Objective Evolutionary Algorithm (AB-MOEA) [21]. The first step of a cycle in AB-MOEA consists in generating new candidates by minimizing the POVs thanks to RVEA. During the second step, the new candidates are re-evaluated by an adaptive function that favors convergence at the beginning and reinforces exploitation as the execution progresses by minimizing the predictive uncertainty delivered by the surrogate. At the third step, q candidates are retained based on an adaptive sampling criterion similar to the reference vector guided replacement of RVEA to promote diversity.

The fourth algorithm is the Surrogate-Assisted Evolutionary Algorithm for Medium Scale Expensive problems (SAEA-ME) [22]. In SAEA-ME, NSGA-II is used to optimize a six-objective acquisition function where the three first objectives are the POVs and the last three objectives are the POVs minus the predictive variances. From the set of proposed candidates, the q ones showing the best hyper-volume improvement considering both the POVs alone and the POVs minus two variances are retained for parallel simulations. SAEA-ME performs well on problems with more than 10 decision variables. The dimensionality reduction feature proposed in [22] is not considered here as it consumes computational budget and can be applied to any method.

A Multi-Task Gaussian Process (MTGP) surrogate model [23] is implemented *via* the GPyTorch library [24] and incorporated into both AB-MOEA and SAEA-ME. Using a MTGP to model multiple objectives has been realized in [25] to control quality in sheet metal forming. In a traditional regression GP [26], a kernel function is specified to model the covariance between the inputs, thus allowing the model to learn the input-output mapping and to return predictions and predictive uncertainties. In the MTGP, inter-task dependencies are also

taken into account in the hope of improving over the case where the tasks are decoupled.

In the present investigation, the tasks are the three objectives and five kernel functions are considered for comparison. The widely used Radial Basis Functions kernel, denoted *rbf* and described in [26], provides very smooth predictors. According to [27], the Matern kernel with hyper-parameter $\nu = 1.5$ or 2.5, called *matern1.5* and *2.5* respectively, is to be preferred to model many physical phenomena. The higher predictive capacity Spectral Mixture kernel proposed in [28] is also raised with 2 and 4 components, denominated *sm2* and *sm4* respectively.

4 Experiments

The computational budget is set to two hours on 18 computing cores, thus allowing 18 simulations to be realized in parallel. The simulation duration varies from one solution to another from 13 to 142 seconds on one computing core. The four competing algorithms are implemented using our pySBO Python tool publicly available at: `https://github.com/GuillaumeBriffoteaux/pySBO`. Ten repetitions of the searches are carried out to ensure statistical robustness of the comparisons. The reference point for hyper-volume calculation is set to an upper bound for each objective $(32.10^6; 32.10^6; 1.5)$.

The surrogate-free approaches NSGA-II, RVEA and RVEA* are equipped with either the *distrib-X*, the *2-points* or the *intermediate* cross-over operator. For NSGA-II, the population size p_s is set to 108 or 162, thus avoiding the idling of the computing cores. For RVEA and RVEA*, we choose $p_s = 105$ or 171 to comply with the constraint imposed by the reference vectors initialization and to keep values close to those imposed for NSGA-II. Ten initial populations composed of 171 simulated solutions are generated to start the algorithms. Each initial population is made at 85% of solutions randomly sampled within the feasible search space and at 15% of candidates picked out on the boundary. When $p_s < 171$, only the best p_s candidates according to the non-dominated sorting defined in [19] are retained. For RVEA and RVEA*, a scaled version of the problem, where the first two objectives are divided by 1000, is also considered to demonstrate the effect of the objectives scales on the behavior of the methods.

The surrogate-based approaches AB-MOEA and SAEA-ME only integrate the *distrib-X* operator and use all the 171 initial samples as initial database. For RVEA in AB-MOEA, $p_s = 105$ and the number of generations is fixed to 20 as recommended in [21], while $p_s = 76$ for NSGA-II in SAEA-ME according to the guidance provided in [22] and the population evolved for 100 generations.

Table 1 shows the ranking of the algorithms according to the final hyper-volumes averaged over the ten repetitions. It can be observed in Table 1 that all the surrogate-based strategies outperform all those without surrogate. In particular, SAEA-ME with the *matern1.5* kernel is the best approach. The MTGP equipped with the *matern1.5* covariance function is preferred in both the SAEA-ME and AB-MOEA frameworks. Regarding the surrogate-free meth-

ods, NSGA-II with the *distrib-X* cross-over mechanism and $p_s = 108$ yields the best averaged hyper-volume. It is worth noticing that the *distrib-X* operator, specifically designed for the problem at hand, is to put forward as it surpasses both the *intermediate* and the *2-points* strategies in all contexts. Among the RVEAs, the best variant is RVEA* with $p_s = 105$ and the *distrib-X* cross-over thus indicating a possibly degenerated or disconnected PF. Indeed, the PF is certainly degenerated as indicates Figure 1 where are plotted the objective vectors from the ten final NDFs obtained by SAEA-ME with the *matern1.5* kernel. When analyzing the influence of objectives scales over the efficiency of RVEA and RVEA*, the conclusions drawn in [20] are confirmed as both algorithms are more appropriate when objectives have similar scales. Indeed, the three objectives lie in $[1655; 13,762]$, $[843; 10,962]$ and $[0; 1]$, respectively. The previous ranges are approximated *a posteriori* based on 250,664 simulations performed in RVEA and RVEA* on the original problem. The necessity to adequately scale the objectives brings a disadvantage to RVEAs as the scaling weights are tedious to define especially in the context of black-box expensive simulations. Another drawback is the constraints on the population size preventing to totally impede the idling of computing cores in all scenarios.

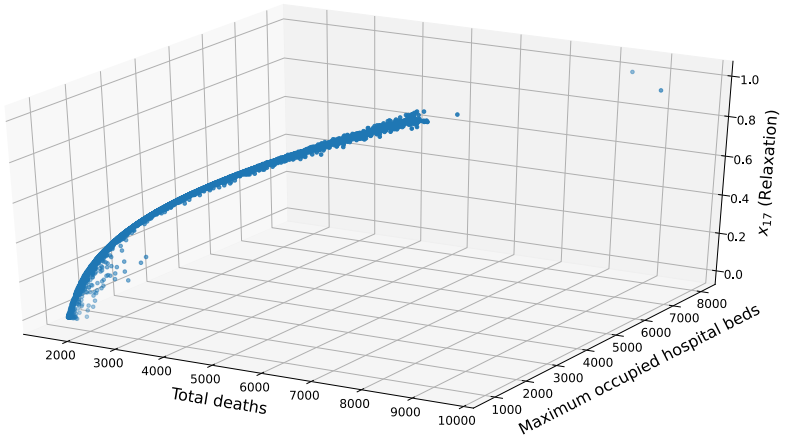

Fig. 1. Best NDFs from the 10 repetitions for SAEA-ME with matern1.5 kernel.

Figure 2 monitors the averaged hyper-volume as the search proceeds for the best strategy per category according to Table 1. The hyper-volume improves sharply at the very beginning of the search for the surrogate-based methods and reaches convergence rapidly (around 300 to 500 simulations). NSGA-II converges much slower but seems not to have converged at the end of the execution. By the right extremities of the curves, it could be expected that the hyper-volume returned by NSGA-II exceeds the one from AB-MOEA for larger numbers of simulations. However, reiterating the experiments for a time budget of four hours has

Algorithm	Cross-over operator	Population size	GP kernel	Objectives scaling	Averaged final Hyper-volume ($\times 10^{10} + 1.535 \times 10^{15}$)
SAEA-ME	**distribX**	**76**	**matern1.5**	-	**80.1800**
SAEA-ME	distribX	76	matern2.5	-	80.1610
SAEA-ME	distribX	76	rbf	-	79.9541
SAEA-ME	distribX	76	sm2	-	79.6701
AB-MOEA	distribX	105	matern1.5	-	79.6200
AB-MOEA	distribX	105	matern2.5	-	79.5879
SAEA-ME	distribX	76	sm4	-	79.5789
AB-MOEA	distribX	105	sm4	-	79.4861
AB-MOEA	distribX	105	sm2	-	79.4841
AB-MOEA	distribX	105	rbf	-	79.4304
NSGA-II	distribX	108	-	-	79.3337
NSGA-II	distribX	162	-	-	79.1876
RVEA*	distribX	105	-	yes	77.2805
RVEA*	distribX	171	-	yes	77.2514
RVEA	distribX	171	-	yes	77.1287
RVEA	distribX	105	-	yes	77.0117
NSGA-II	intermediate	108	-	-	76.9946
NSGA-II	intermediate	162	-	-	76.8320
NSGA-II	2-points	162	-	-	75.6959
NSGA-II	2-points	108	-	-	75.5889
RVEA	distribX	105	-	-	75.5184
RVEA*	intermediate	171	-	yes	75.3816
RVEA*	intermediate	105	-	yes	75.2841
RVEA	intermediate	171	-	yes	75.2006
RVEA	intermediate	105	-	yes	75.1562
RVEA*	distribX	105	-	-	75.1555
RVEA*	distribX	171	-	-	75.1372
RVEA	distribX	171	-	-	75.0563
RVEA*	2-points	171	-	yes	74.9803
RVEA	2-points	171	-	yes	74.9195
RVEA	2-points	105	-	yes	74.7692
RVEA*	2-points	105	-	yes	74.7535
RVEA*	intermediate	105	-	-	74.5607
RVEA	intermediate	105	-	-	74.5585
RVEA	intermediate	171	-	-	74.4959
RVEA*	intermediate	171	-	-	74.4266
RVEA	2-points	171	-	-	74.3694
RVEA	2-points	105	-	-	74.3518
RVEA*	2-points	171	-	-	74.3264
RVEA*	2-points	105	-	-	74.2507

Table 1. Ranking of the surrogate-based and surrogate-free approaches according to the averaged final hyper-volumes over the 10 repetitions.

not allowed to verify this expectation. Figure 2 specifies that the impact of objectives scaling on RVEAs appears from around 300 simulations. In the setting of a capped computational budget, it is important to strongly favor convergence and exploitation at the onset of the search. SAEA-ME and AB-MOEA realizes this by minimizing the POVs at the top beginning of the execution. The difference between the two approaches lies in the incorporation of the predictive uncertainty. In SAEA-ME, a degree of exploration is maintained by maximization of the predictive variance. Conversely, minimization of the predictive uncertainty is involved at latter stages in AB-MOEA. In spite of the convergence-oriented strategy adopted by RVEAs at the early stages of the search, the embedded mechanism set up to handle many objectives is quite heavy and reveals to be unsuitable when the computational budget is restricted. Indeed, in [20] the algorithms are run from 500 to 1,000 generations while 10 to 20 generations are allowed by our computational budget.

Reducing the solving time of moderately expensive optimization problems where the simulation lasts less than five minutes may enable to manage optimization under uncertainty. As the calibration of the simulation tool is uncertain, multiple configurations of its parameters can be considered, resulting in multiple optimization exercises to be executed and thus enabling to gain insight about the variability of the results.

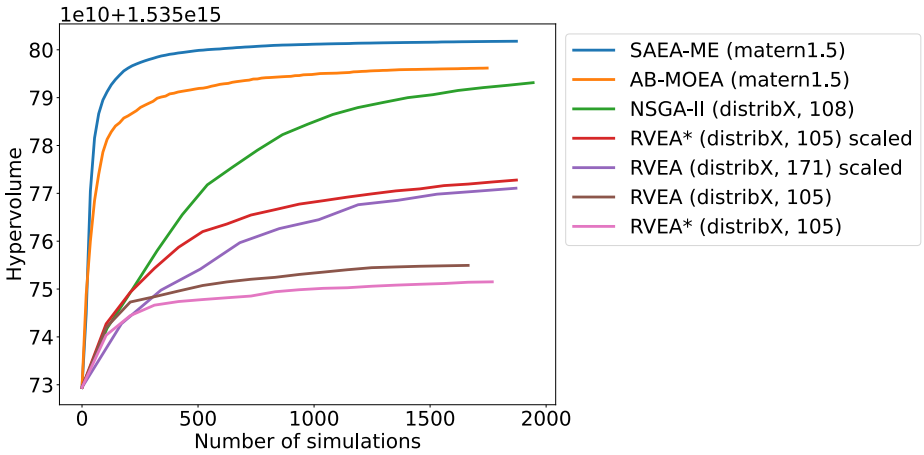

Fig. 2. Averaged hyper-volume according to the number of simulations.

The optimal allocation plan implies providing 70% of the doses to the 10-19 years old age-group and 30% to the 20-29 age-group during phase 2 according to Figure 3. In phase 3, 70% of the doses are assigned to 20-29 years old individuals and 15% to both the 40-49 and 10-19 age-categories. This plan prioritizes the vaccination of younger adults as they are the most transmitting cohort because of their high contact rate in the population [29]. Nevertheless, the present re-

sults have to be considered with caution. Since our experiments date back to the beginning of 2021, few feedback about vaccination efficiency was available. It is assumed here that the vaccine reduces transmission although it might not be the case for the Omicron variant of concern that started to break through the world at the end of 2021. Our results are similar to those presented in [30, 31] for influenza. From Figure 4 where the total number of deaths and the maximum number of occupied hospital beds are displayed with respect to the relaxation variable x_{17}, the alleviation of the physical distancing reveals to trigger an augmentation of the hospital occupancy and deaths.

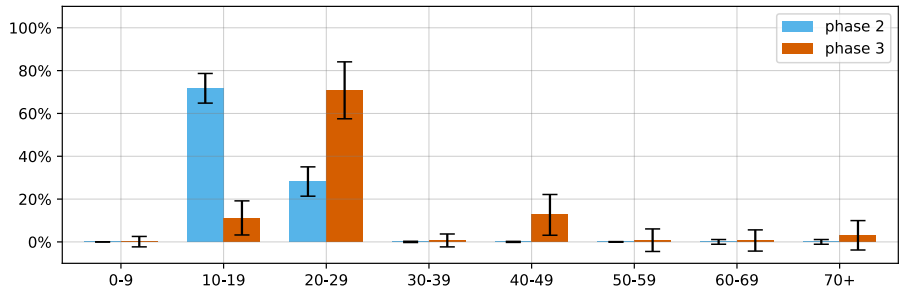

Fig. 3. Vaccines distribution according to age-categories. Averaged solutions from the best final NDFs returned by the 10 repetitions for SAEA-ME with matern1.5 kernel.

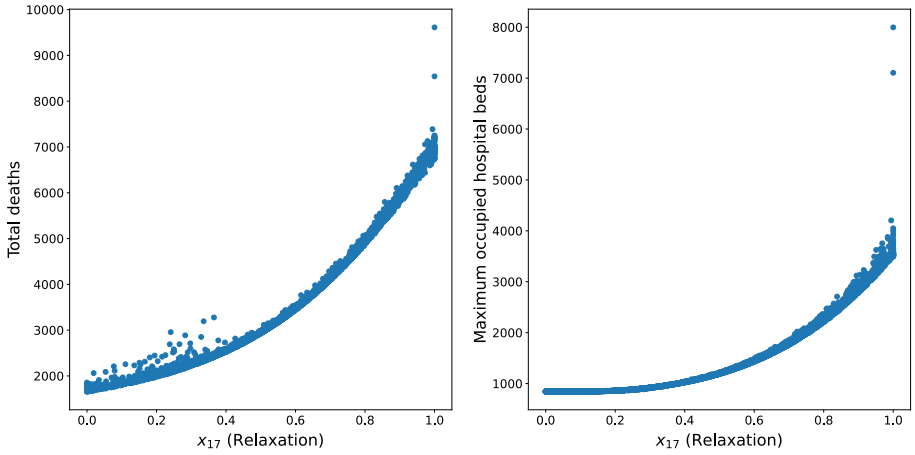

Fig. 4. Total number of deaths and maximum number of occupied hospital beds according to relaxation of the physical distancing x_{17}. Best NDFs from the 10 repetitions for SAEA-ME with matern1.5 kernel.

5 Conclusion

This paper demonstrates the suitability of parallel surrogate-based multi-objective optimization algorithms to handle the moderately computationally expensive COVID-19 vaccines allocation problem for Malaysia. In particular, SAEA-ME provides reliable results in a fast way. As future works, we suggest to benefit from the computational cost reduction of black-box simulation-based problem solving to take the uncertainty around the calibration of the simulator into account.

References

1. S. Chang *et al.* Modelling transmission and control of the covid-19 pandemic in australia. *Nature Communications*, 11,5710, 03 2020.
2. J. M. Trauer, M. J. Lydeamore, G. W. Dalton, D. Pilcher, M. T. Meehan, E. S. McBryde, A. C. Cheng, B. Sutton, and R. Ragonnet. Understanding how victoria, australia gained control of its second covid-19 wave. *Nature Communications*, 12(6266), 2021.
3. D. Duque, D. P. Morton, B. Singh, Z. Du, R. Pasco, and L. A. Meyers. Timing social distancing to avert unmanageable covid-19 hospital surges. *Proceedings of the National Academy of Sciences*, 117(33):19873–19878, 2020.
4. L. Matrajt, J. Eaton, T. Leung, D. Dimitrov, J. T. Schiffer, D. A. Swan, and H. Janes. Optimizing vaccine allocation for covid-19 vaccines: potential role of single-dose vaccination. *Nature Communications*, 12(3449), 2021.
5. L. Matrajt and I. Longini. Optimizing vaccine allocation at different points in time during an epidemic. *PloS one*, 5:e13767, 11 2010.
6. L. Matrajt, J. Eaton, T. Leung, and E. R. Brown. Vaccine optimization for covid-19: Who to vaccinate first? *Science Advances*, 7(6):eabf1374, 2021.
7. C. Buhat *et al.* Using constrained optimization for the allocation of covid-19 vaccines in the philippines. *Applied health economics and health policy*, 19(5):699–708, 2021.
8. S. Han *et al.* Time-varying optimization of covid-19 vaccine prioritization in the context of limited vaccination capacity. *Nature communications*, 12(1):4673, August 2021.
9. H. Anahideh, L. Kang, and N. Nezami. Fair and diverse allocation of scarce resources. *Socio-Economic Planning Sciences*, page 101193, 2021.
10. K. M. Bubar, K. Reinholt, S. M. Kissler, M. Lipsitch, S. Cobey, Y. H. Grad, and D. B. Larremore. Model-informed covid-19 vaccine prioritization strategies by age and serostatus. *Science*, 371(6532):916–921, 2021.
11. E. S. McBryde, M. T. Meehan, J. M. Caldwell, A. I. Adekunle, S. T. Ogunlade, M. A. Kuddus, R. Ragonnet, P. Jayasundara, J. M. Trauer, and R. C. Cope. Modelling direct and herd protection effects of vaccination against the sars-cov-2 delta variant in australia. *Medical Journal of Australia*, 215(9):427–432, 2021.
12. G. Briffoteaux, M. Gobert, R. Ragonnet, J. Gmys, M. Mezmaz, N. Melab, and D. Tuyttens. Parallel surrogate-assisted optimization: Batched bayesian neural network-assisted ga versus q-ego. *Swarm and Evolutionary Computation*, 57:100717, 2020.
13. J. M. Cicchese, E. Pienaar, D. E. Kirschner, and J. J. Linderman. Applying optimization algorithms to tuberculosis antibiotic treatment regimens. *Cellular and Molecular Bioengineering*, 10,6:523–535, December 2017.

14. R. Miikkulainen *et al.* From prediction to prescription: Evolutionary optimization of nonpharmaceutical interventions in the covid-19 pandemic. *IEEE Transactions on Evolutionary Computation*, 25(2):386–401, 2021.
15. T. N. Vilches *et al.* Covid-19 hospitalizations and deaths averted under an accelerated vaccination program in northeastern and southern regions of the usa. *The Lancet Regional Health - Americas*, 6:100147, 2022.
16. M. Sheel, S. McEwen, and S. E. Davies. Brand inequity in access to covid-19 vaccines. *The Lancet Regional Health - Western Pacific*, 18:100366, 2022.
17. E. G. Talbi. *Metaheuristics: From Design to Implementation.* Wiley Series on Parallel and Distributed Computing. Wiley, 2009.
18. Z. Michalewicz, D. Dasgupta, R. G. Le Riche, and M. Schoenauer. Evolutionary algorithms for constrained engineering problems. *Computers and Industrial Engineering*, 30(4):851–870, 1996.
19. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
20. R. Cheng, Y. Jin, M. Olhofer, and B. Sendhoff. A reference vector guided evolutionary algorithm for many-objective optimization. *IEEE Transactions on Evolutionary Computation*, 20(5):773–791, 2016.
21. X. Wang, Y. Jin, S. Schmitt, and M. Olhofer. An adaptive bayesian approach to surrogate-assisted evolutionary multi-objective optimization. *Information Sciences*, 519:317–331, 2020.
22. X. Ruan, K. Li, B. Derbel, and A. Liefooghe. Surrogate assisted evolutionary algorithm for medium scale multi-objective optimisation problems. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, GECCO '20, page 560–568, New York, NY, USA, 2020. Association for Computing Machinery.
23. E. V. Bonilla, K. Chai, and C. Williams. Multi-task gaussian process prediction. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2008.
24. J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*, 2018.
25. W. Xia, Huan Yang, Xiaoping Liao, and Jiangbang Zeng. A multi-objective optimization method based on gaussian process simultaneous modeling for quality control in sheet metal forming. *The International Journal of Advanced Manufacturing Technology*, 72:1333–1346, 2014.
26. C. E. Rasmussen. *Gaussian processes for machine learning.* MIT Press, 2006.
27. M.L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging.* Springer Series in Statistics. Springer New York, 1999.
28. Andrew Gordon Wilson and Ryan Prescott Adams. Gaussian process kernels for pattern discovery and extrapolation, 2013.
29. K. Prem, A. R. Cook, and M. Jit. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLOS Computational Biology*, 13(9):1–21, 09 2017.
30. D. Weycker, J. Edelsberg, M. E. Halloran, I. M. Longini, A. Nizam, V. Ciuryla, and G. Oster. Population-wide benefits of routine vaccination of children against influenza. *Vaccine*, 23(10):1284–1293, 2005.
31. J. Medlock and A. P. Galvani. Optimizing influenza vaccine distribution. *Science*, 325(5948):1705–1708, 2009.