

Towards an Evolvable Cancer Treatment Simulator

Richard J. Preen^{a,*}, Larry Bull^a, Andrew Adamatzky^b

^a*Department of Computer Science and Creative Technologies*

^b*Unconventional Computing Laboratory*

University of the West of England, Bristol, BS16 1QY, UK

Abstract

The use of high-fidelity computational simulations promises to enable high-throughput hypothesis testing and optimisation of cancer therapies. However, increasing realism comes at the cost of increasing computational requirements. This article explores the use of surrogate-assisted evolutionary algorithms to optimise the targeted delivery of a therapeutic compound to cancerous tumour cells with the multicellular simulator, PhysiCell. The use of both Gaussian process models and multi-layer perceptron neural network surrogate models are investigated. We find that evolutionary algorithms are able to effectively explore the parameter space of biophysical properties within the agent-based simulations, minimising the resulting number of cancerous cells after a period of simulated treatment. Both model-assisted algorithms are found to outperform a standard evolutionary algorithm, demonstrating their ability to perform a more effective search within the very small evaluation budget. This represents the first use of efficient evolutionary algorithms within a high-throughput multicellular computing approach to find therapeutic design optima that maximise tumour regression.

Keywords: Agent-based model, evolutionary algorithm, cancer, PhysiCell, high-throughput computing, surrogate modelling

1. Introduction

PhysiCell ([Ghaffarizadeh et al., 2018](#)) is an open source¹ multicellular simulator based on the biotransport solver, BioFVM ([Ghaffarizadeh et al., 2016](#)). BioFVM simulates

*Corresponding author

Email address: richard2.preen@uwe.ac.uk (Richard J. Preen)

¹PhysiCell source code (BSD license): <https://github.com/MathCancer/PhysiCell>

Preprint submitted to *BioSystems*

May 10, 2019

substrate secretion, diffusion, uptake, and decay; while PhysiCell models the emergent tissue-scale behaviour by simulating cell cycling, death, volume, mechanics, and motility. PhysiCell enables the simulation of new environmental substrates, cell types, and systems of cells, providing a general-purpose toolkit for exploring multicellular systems.

[Norton et al. \(2019\)](#) present a review of agent-based models used to study cancer-immune interaction and immunotherapy; and [Metzcar et al. \(2019\)](#) provide a more general overview of cell-based computational modelling in cancer biology. Many proposed cancer therapies attempt to target malignant cells by finding specific surfaces or molecules in order that drugs can be conjugated to custom antibodies or encapsulated in custom nanoparticles. [Ghaffarizadeh et al. \(2018\)](#) introduced the agent-based PhysiCell 2-D anti-cancer biorobots simulation. This began the design of cell-cell interaction rules to create a multicellular cargo delivery system that actively delivers a cancer therapeutic beyond regular drug transport limits to hypoxic cancer regions. These model rules were manually tuned to achieve this (as yet unoptimised) design objective, requiring weeks of people-hours to configure, code, test, visualise, and evaluate ([Ozik et al., 2018](#)).

[Ghaffarizadeh et al. \(2018\)](#) also presented 3-D simulations of cancer immunotherapy. Using this simulator, [Ozik et al. \(2018\)](#) performed a human-selected parameter sweep (27 parameter sets; each set sampled 10 times) with each simulation requiring ≈ 2 days to complete. The results provided insights into therapeutic failure, thus demonstrating the potential of high-throughput computing to investigate high dimensional cancer simulator parameter spaces. High-throughput model investigation and hypothesis testing provides a new paradigm for solving complex problems, gaining new insights, and improving cancer treatment strategies ([Ozik et al., 2018](#)).

Surrogate model-based optimisation has long been used in applications requiring expensive parameter evaluations, whether via simulated or physical testing ([Sacks et al., 1989](#); [Jones et al., 1998](#); [Settles, 2010](#)). In this article, we explore the use of surrogate-assisted evolutionary algorithms (EAs) to sequentially optimise the targeted delivery of a therapeutic compound to cancerous tumour cells with the multicellular simulator, PhysiCell. This represents the first use of efficient EAs within a high-throughput multicellular computing approach to find therapeutic design optima that maximise tumour regression.

2. Background

The use of surrogate models to reduce the number of costly EA fitness evaluations can be traced back to the origins of the discipline (Dunham et al., 1963). Jin et al. (2018) present an overview of the process: first the variables to optimise are chosen; some initial parameters are then evaluated; a surrogate (regression) model is selected and used to build a model of the evaluated parameters; followed by a search of the model to identify new parameter values to evaluate; and the newly evaluated parameters added to the existing data. The process loops via continued iterations of model building, searching, and evaluation of selected parameters.

Many regression models have been used as a surrogate for the real fitness function, including multi-layer perceptron (MLP) based artificial neural networks (e.g., Bull, 1997), Gaussian processes (GP; also known as Kriging; e.g., Liu et al., 2014), radial basis functions (RBF; e.g., Regis, 2014), support vector regression (e.g., Yun et al., 2009), particle swarm optimisation (e.g., Wang et al., 2017), Markov networks (e.g., Brownlee et al., 2013), and coevolved fitness predictors (Schmidt and Lipson, 2008). Many ensembles have also been explored (e.g., Wang et al., 2018). Probabilistic models such as GP are perhaps the most widely used surrogate since they provide a measure of confidence that can be used to efficiently select samples for evaluation (Jin et al., 2018). Preen and Bull (2016) showed that with very small and noisy samples there is little difference between the modelling approaches, with MLPs appearing to be a robust approach capable of capturing the underlying structure of the search space.

The computational complexity of model building relative to the sampling expense is a key consideration, particularly in the case where large archive sets are used for training; for example, the computation time for GP training increases in cubic with the number of training data (Shahriari et al., 2016). MLP models are typically adopted when there are a large number of decision variables and/or training data (Chugh et al., 2017). This leads to a further key consideration, which is whether to use the full archive set for model training. It may be necessary to restrict the number of samples for use with a computationally complex model such as GP when there is a large archive set (Jin et al., 2018). Additionally, if there are any significant temporal affects, issues surround how best to select the subset for training (Preen and Bull, 2017).

The task of selecting which samples to evaluate is analogous with the problem of active learning wherein the algorithm is able to interactively query an oracle (e.g., a user) to obtain the output at a given data point (Settles, 2010). For online data-driven optimisation, initial data collection is often simply performed via random selection or a design of experiments technique such as Latin hypercube sampling (Jin et al., 2018). Subsequent samples are selected through the use of an acquisition function (also known as the infill sampling criteria) which rates the expected utility of evaluating a candidate solution (Shahriari et al., 2016). The most commonly used acquisition functions are the mean of model prediction, and the expected improvement (EI; Jiao et al., 2019). Acquisition functions typically aggregate the model predicted fitness (e.g., mean) and estimated confidence (e.g., standard deviation) to explore regions of both high promise and high uncertainty. In addition to variance-based sampling, query-by-committee, cross-validation, and gradient-based methods also exist (Liu et al., 2018).

The use of the acquisition function is performed within an overall model management. Jin (2011) categorised approaches to model management as generational, individual, or population-based, depending upon whether whole generations, samples of individuals within a generation, or sub-populations are evaluated. The pre-selection approach (Emmerich et al., 2002) uses the model to rate M number of offspring and the best of these are chosen for evaluation on the real fitness function. As highlighted by Jin (2011), the main difference between individual-based strategies and pre-selection is that the real fitness value is always used for selection, whereas in individual-based methods selection may be based on fitness values from the surrogate. This potentially makes the approach more robust to noise. Pre-selection has previously been shown effective for the evolution of (noisy) physical systems in which each design is physically instantiated and evaluated for fitness (Preen and Bull, 2015, 2016, 2017).

Four main categories of uncertainty (or noise) are frequently encountered in real-world optimisation problems (Jin and Branke, 2005):

1. Uncertainty in the objective function evaluation. This is common in stochastic simulations and physical systems where evaluations with the same parameters produce different results, e.g., by sensory measurement errors. Efficiently reducing the objective function error requires optimally selecting how many samples to perform

for a given parameter set. Issues such as how best to aggregate the results, e.g., by mean or median may also affect convergence.

2. The objective function is approximated by a surrogate model. Efficiently reducing the model error requires optimally selecting which samples to evaluate with the real objective function to provide new training data and avoidance of over-fitting through validation (Bischl et al., 2012). Approximated objective values are also encountered where the function evaluations take place over time and partial fitness scores can be observed by early termination (Park et al., 2013). Approximation error can be beneficial to the search process by aiding the escape of local optima (Lim et al., 2010).
3. Dynamic optimisation wherein the objectives change over time (e.g., Chen et al., 2018). In these scenarios, population diversity and memory mechanisms are essential.
4. Robust optimisation (e.g., Yu and Suprayitno, 2017) which seeks solutions that are less sensitive to small parameter perturbations, e.g., due to manufacturing tolerances. This typically involves a trade-off between solution quality and robustness.

Rakshit et al. (2017) present an overview on techniques for dealing with noisy evolutionary optimisation. Uncertainty in the objective function can be addressed via the explicit averaging of resampled parameters or implicit averaging using a large population size. The evolutionary selection mechanism may also be modified to account for noisy evaluations; for example, by only accepting offspring with observed fitness greater than the parent’s plus some threshold. Repeated sampling is typically more effective for noise handling than parent populations and threshold selection, resulting in an exponential speed-up for noisy evolutionary optimisation in some cases (Qian et al., 2018). The number of samples to perform for a given parameter set can either remain fixed (static) for all candidate solutions or be dynamically allocated to each, e.g., based on the sample variance (Siegmund et al., 2013).

3. Physics-based Multicellular Simulations

3.1. Methodology

Here we use the PhysiCell 2-D anti-cancer biorobots simulator ([Ghaffarizadeh et al., 2018](#)). This performs a multicellular simulation of targeted drug delivery by modelling three types of cells:

- Oxygen consuming cancer cells that generate a chemoattractant c_1 forming an oxygen gradient which can be used to guide worker cells.
- Worker cells that may adhere to cargo cells. Each cell has a persistence time, migration speed, migration direction, and migration bias. Worker cells perform a biased random migration towards cancer cells when adhered to cargo, and a biased random migration towards cargo cells when unadhered. Migration biases range $[0,1]$ with 0 representing Brownian motion and 1 deterministic motion. The motility of unadhered worker cells is disabled if c_1 falls below a threshold.
- Cargo cells that secrete a diffusible chemoattractant c_2 used to guide worker cells. Adhered cargo cells detach from worker cells and secrete a therapeutic compound that induces apoptosis in nearby tumour cells when oxygenation falls below the cargo release o_2 threshold.

Each simulation is initialised with a 200 micron radius tumour. After 7 simulated days of tumour growth, 500 cells are “injected” near the tumour: 10% worker cells and 90% cargo cells. The simulation subsequently continues for 3 additional days of cancer therapy. A single simulation requires ≈ 5 minutes of wall-clock time on an Intel[®] Xeon[®] CPU E5-2650 v4 @ 2.20GHz with 64GB RAM using half of the 48 cores. Here we search the $N = 6$ parameters specifying worker agent characteristics and cargo properties with the goal of minimising the resulting number of cancer cells after a period of simulated treatment. All other parameters remain at their original values as shown in Table 1.

In particular, we explore the attached worker migration bias $[0,1]$; the unattached worker migration bias $[0,1]$; worker relative adhesion $[0,10]$; worker relative repulsion $[0,10]$; worker motility persistence time (minutes) $[0,10]$; and the cargo release o_2 threshold (mmHg) $[0,20]$. We adopt the surrogate-assisted pre-selection approach (and pa-

Table 1: Default multicellular simulation parameters.

Maximum attachment distance	18 microns
Minimum attachment distance	14 microns
Worker apoptosis rate	0 minutes ⁻¹
Worker migration speed	2 microns/minute
Worker o ₂ relative uptake	0.1 minutes ⁻¹
Cargo o ₂ relative uptake	0.1 minutes ⁻¹
Cargo apoptosis rate	4.065e-5 minutes ⁻¹
Cargo relative adhesion	0
Cargo relative repulsion	5
Damage rate	0.03333 minutes ⁻¹
Repair rate	0.004167 minutes ⁻¹
Drug death rate	0.004167 minutes ⁻¹
Maximum relative cell adhesion distance	1.25
Elastic coefficient	0.05 minutes ⁻¹
Maximum elastic displacement	50 microns
Motility shutdown detection threshold	0.001
Attachment receptor threshold	0.1

parameter values) previously used successfully to perform physical test-driven optimisation (Preen and Bull, 2017). As benchmark, a steady-state genetic algorithm (GA) with population size $P = 20$ is used; tournament size $T = 3$ for both selection and replacement; uniform crossover is performed with $\mathcal{X} = 80\%$ probability; and a per allele mutation rate $\mu = 1/N$ with a uniform random step size $s = [-5, 5]\%$. A static sampling approach is used wherein k simulation runs are performed for each candidate solution before assigning the fitness as the mean number of remaining cells after 7 simulated days of tumour growth plus 3 days of targeted drug delivery.

For the surrogate-assisted GA, all individuals in the initial population are evaluated and a regression model fit. Subsequently, evolution proceeds by iteratively selecting 2 parents via tournament and then creating and evaluating $M = 1000$ offspring with the model via an acquisition function. The most promising of these is then selected for

evaluation by the multicellular simulator and replaces an individual in the population selected via a negative tournament. Finally, the evaluated archive set is updated and the model retrained. The full archive set is used for training since there are no temporal affects on sampling, the time required to fit the model is insignificant in comparison with the sample evaluation time, and only a few decision variables are optimised. Algorithm 1 provides an outline of the surrogate-assisted GA with static sampling.

Since MLPs are extrapolative and GPs are interpolative, here we employ both approaches to model building and observe their effects on the evolved biophysical parameters. Here, the GP model (Rasmussen and Williams, 2006) uses an RBF kernel function and the MLP uses $H = 10$ rectified linear units in the hidden layer. Both models are trained using the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton optimisation algorithm (Byrd et al., 1995). As the GP model provides a measure of confidence (i.e., standard deviation) for each prediction, the EI is used as the acquisition (rating) function to select the next offspring; whereas the MLP does not provide a measure of confidence and so the best predicted fitness is used in this case. Algorithm 2 shows the EI rating function used for the GP model. For the MLP model, *Rating()* in Algorithm 1 returns the model predicted fitness. All experiments are initialised with the same randomly generated population.

3.2. Results

Since each simulation run is costly, we initially explored the case with $k = 1$. However, after 400 evaluations/simulations, the performance of the fittest individual discovered by the GA was not significantly different than the fittest individual in the initial population, $p > 0.05$ using a Wilcoxon rank-sums test, showing that the GA is severely misled by the significant variance in simulation runs with the same parameter set.

Figure 1 shows the mean number of cells resulting from the fittest individual discovered for each algorithm with $k = 10$. After evaluating 200 candidates (2000 simulations), the best GP-assisted solution (mean = 852.10, SD = 40.43, samples = 10, min = 773, median = 870, kurtosis = -0.89) and the best MLP-assisted solution (mean = 852.20, SD = 39.84, samples = 10, min = 791, median = 853, kurtosis = -0.74) are significantly less than the best GA solution without surrogate assistance (mean = 889.30, SD = 32.75, samples = 10, min = 831, median = 898, kurtosis = -0.83), $p \leq 0.05$ using a Wilcoxon

rank-sums test. There is no significant difference between the surrogate models. All algorithms found solutions with a significantly lower mean number of cells than the best individual in the initial population (mean = 953.50, SD = 41.59, samples = 10, min = 887, median = 947, kurtosis = -1.10), $p \leq 0.05$.

Figure 2 shows scatter plots of all evaluated individuals for each evolved parameter. As can be seen, both surrogate-assisted EAs identified best solutions with a mean of 852 cells and these solutions have similar values for 4 of the parameters: ≈ 0.5 unattached worker migration bias; ≈ 6 worker relative adhesion; ≈ 10 worker motility persistence time; and ≈ 11 cargo release O_2 threshold. However, the MLP-assisted model achieved this with an attached worker migration bias of 0.89, whereas the GP-assisted solution was 0.29. Additionally, the worker relative repulsion was 5.9, compared with 1.13 for the GP-assisted model.

There appears to be a clear funnel with a minimum at ≈ 11 for the cargo release O_2 threshold, suggesting that this is the global optima for the parameter. Ghaffarizadeh et al. (2018) used an initial cargo release O_2 threshold of 10, finding that “once enough cancer cells were killed, hypoxia was reduced so that worker cells clustered near the oxygen minimum, but no longer released their cargo”. Increasing the threshold to 15 “reduced but did not eliminate this behaviour”. The results of these simulations suggest that using a threshold of 11 results in the best performance. As Ghaffarizadeh et al. (2018) note, “the cargo release rules need to be carefully engineered. Such a system could potentially activate and deactivate to keep a tumour cell population in control, and to reduce hypoxia [which is known to drive cancer cell adaptation to more aggressive phenotypes (Wilson and Hay, 2011; Eales et al., 2016)]”.

An example run of the fittest evolved individual with the GP-surrogate model is shown in Figure 3, showing that the worker cells appear to disperse evenly and effectively deliver the cargo to the tumour.

4. Conclusions

This article has shown that EAs are able to effectively explore the parameter space of biophysical properties within the agent-based multicellular simulator, PhysiCell. EAs successfully minimised the number of cancerous cells after a period of simulated treat-

ment. Both surrogate-assisted algorithms were found to outperform the standard GA, thereby reducing the number of expensive simulations required. No significant difference in the resulting number of cancerous cells was observed between the models, showing the robustness of the overall pre-selection with static sampling approach employed. By exploring both surrogate models that extrapolate and models that interpolate, we were able to identify different parameter sets that achieved similar reductions in cancerous cells. Thus, we have demonstrated the use of efficient EAs within a high-throughput computing approach to find therapeutic design optima that maximise tumour regression.

Acknowledgement

This work was supported by the European Research Council under the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 800983.

References

- Bischl, B., Mersmann, O., Trautmann, H., Weihs, C., 2012. Resampling methods for meta-model validation with recommendations for evolutionary computation. *Evol. Comput.* 20, 249–275. doi:[10.1162/EVC0_a_00069](https://doi.org/10.1162/EVC0_a_00069).
- Brownlee, A.E.I., McCall, J.A.W., Zhang, Q., 2013. Fitness modeling with Markov networks. *IEEE Trans. Evol. Comput.* 17, 862–879. doi:[10.1109/TEVC.2013.2281538](https://doi.org/10.1109/TEVC.2013.2281538).
- Bull, L., 1997. Model-based evolutionary computing: A neural network and genetic algorithm architecture, in: *Proc. IEEE Int. Conf. Evol. Comput.*, pp. 611–616. doi:[10.1109/ICEC.1997.592384](https://doi.org/10.1109/ICEC.1997.592384).
- Byrd, R.H., Lu, P., Nocedal, J., Zhu, C., 1995. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* 16, 1190–1208. doi:[10.1137/0916069](https://doi.org/10.1137/0916069).
- Chen, R., Li, K., Yao, X., 2018. Dynamic multiobjectives optimization with a changing number of objectives. *IEEE Trans. Evol. Comput.* 22, 157–171. doi:[10.1109/TEVC.2017.2669638](https://doi.org/10.1109/TEVC.2017.2669638).
- Chugh, T., Sindhya, K., Hakanen, J., Miettinen, K., 2017. A survey on handling computationally expensive multiobjective optimization problems with evolutionary algorithms. *Soft Comput.* doi:[10.1007/s00500-017-2965-0](https://doi.org/10.1007/s00500-017-2965-0). in press.
- Dunham, B., Fridshal, R., North, J.H., 1963. Design by natural selection. *Synthese* 15, 254–259. doi:[10.1007/BF00484855](https://doi.org/10.1007/BF00484855).
- Eales, K.L., Hollinshead, K.E.R., Tennant, D.A., 2016. Hypoxia and metabolic adaptation of cancer cells. *Oncogenesis* 5, e190. doi:[10.1038/oncsis.2015.50](https://doi.org/10.1038/oncsis.2015.50).
- Emmerich, M., Giotis, A., Özdemir, M., Bäck, T., Giannakoglou, K., 2002. Metamodel-assisted evolution strategies, in: Guervós, J.J.M., Adamidis, P., Beyer, H.G., Schwefel, H.P., Fernández-Villacañas, J.L.

- (Eds.), Proc. PPSN VII. Springer, Berlin, Germany. volume 2439 of *LNCS*, pp. 361–370. doi:[10.1007/3-540-45712-7_35](https://doi.org/10.1007/3-540-45712-7_35).
- Ghaffarizadeh, A., Friedman, S.H., Macklin, P., 2016. BioFVM: An efficient, parallelized diffusive transport solver for 3-D biological simulations. *Bioinformatics* 32, 1256–1258. doi:[10.1093/bioinformatics/btv730](https://doi.org/10.1093/bioinformatics/btv730).
- Ghaffarizadeh, A., Heiland, R., Friedman, S.H., Mumenthaler, S.M., Macklin, P., 2018. PhysiCell: An open source physics-based cell simulator for 3-D multicellular systems. *PLoS Comput. Biol.* 14, 1–31. doi:[10.1371/journal.pcbi.1005991](https://doi.org/10.1371/journal.pcbi.1005991).
- Jiao, R., Zeng, S., Li, C., Jiang, Y., Jin, Y., 2019. A complete expected improvement criterion for Gaussian process assisted highly constrained expensive optimization. *Inf. Sci.* 471, 80–96. doi:[10.1016/j.ins.2018.09.003](https://doi.org/10.1016/j.ins.2018.09.003).
- Jin, Y., 2011. Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm Evol. Comput.* 1, 61–70. doi:[10.1016/j.swevo.2011.05.001](https://doi.org/10.1016/j.swevo.2011.05.001).
- Jin, Y., Branke, J., 2005. Evolutionary optimization in uncertain environments—A survey. *IEEE Trans. Evol. Comput.* 9, 303–317. doi:[10.1109/TEVC.2005.846356](https://doi.org/10.1109/TEVC.2005.846356).
- Jin, Y., Wang, H., Chugh, T., Guo, D., Miettinen, K., 2018. Data-driven evolutionary optimization: An overview and case studies. *IEEE Trans. Evol. Comput.* doi:[10.1109/TEVC.2018.2869001](https://doi.org/10.1109/TEVC.2018.2869001). in press.
- Jones, D.R., Schonlau, M., Welch, W.J., 1998. Efficient global optimization of expensive black-box functions. *J. Global Optim.* 13, 455–492. doi:[10.1023/A:1008306431147](https://doi.org/10.1023/A:1008306431147).
- Lim, D., Jin, Y., Ong, Y.S., Sendhoff, B., 2010. Generalizing surrogate-assisted evolutionary computation. *IEEE Trans. Evol. Comput.* 14, 329–355. doi:[10.1109/TEVC.2009.2027359](https://doi.org/10.1109/TEVC.2009.2027359).
- Liu, B., Zhang, Q., Gielen, G.G.E., 2014. A Gaussian process surrogate model assisted evolutionary algorithm for medium scale expensive optimization problems. *IEEE Trans. Evol. Comput.* 18, 180–192. doi:[10.1109/TEVC.2013.2248012](https://doi.org/10.1109/TEVC.2013.2248012).
- Liu, H., Ong, Y.S., Cai, J., 2018. A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design. *Struct. Multidisc. Optim.* 57, 393–416. doi:[10.1007/s00158-017-1739-8](https://doi.org/10.1007/s00158-017-1739-8).
- Metzcar, J., Wang, Y., Heiland, R., Macklin, P., 2019. A review of cell-based computational modeling in cancer biology. *JCO Clin. Cancer Inform.* 3, 1–13. doi:[10.1200/CCI.18.00069](https://doi.org/10.1200/CCI.18.00069).
- Norton, K.A., Gong, C., Jamalian, S., Popel, A.S., 2019. Multiscale agent-based and hybrid modeling of the tumor immune microenvironment. *Processes* 7, 37. doi:[10.3390/pr7010037](https://doi.org/10.3390/pr7010037).
- Ozik, J., Collier, N.T., Wozniak, J.M., Macal, C.M., Cockrell, C., Friedman, S.H., Ghaffarizadeh, A., Heiland, R., An, G., Macklin, P., 2018. High-throughput cancer hypothesis testing with an integrated PhysiCell-EMEWS workflow. *BMC Bioinform.* 19, 483. doi:[10.1186/s12859-018-2510-x](https://doi.org/10.1186/s12859-018-2510-x).
- Park, N., Kim, K., McKay, R.I., 2013. Cutting evaluation costs: An investigation into early termination in genetic programming, in: Proc. IEEE Congr. Evol. Comput., IEEE Press, Piscataway, NJ, USA. pp. 3291–3298. doi:[10.1109/CEC.2013.6557973](https://doi.org/10.1109/CEC.2013.6557973).
- Preen, R.J., Bull, L., 2015. Toward the coevolution of novel vertical-axis wind turbines. *IEEE Trans. Evol. Comput.* 19, 284–294. doi:[10.1109/TEVC.2014.2316199](https://doi.org/10.1109/TEVC.2014.2316199).

- Preen, R.J., Bull, L., 2016. Design mining interacting wind turbines. *Evol. Comput.* 24, 89–111. doi:[10.1162/EVCO_a_00144](https://doi.org/10.1162/EVCO_a_00144).
- Preen, R.J., Bull, L., 2017. On design mining: Coevolution and surrogate models. *Artif. Life* 23, 186–205. doi:[10.1162/ARTL_a_00225](https://doi.org/10.1162/ARTL_a_00225).
- Qian, C., Yu, Y., Tang, K., Jin, Y., 2018. On the effectiveness of sampling for evolutionary optimization in noisy environments. *Evol. Comput.* 26, 237–267. doi:[10.1162/evco_a_00201](https://doi.org/10.1162/evco_a_00201).
- Rakshit, P., Konar, A., Das, S., 2017. Noisy evolutionary optimization algorithms – A comprehensive survey. *Swarm Evol. Comput.* 33, 18–45. doi:[10.1016/j.swevo.2016.09.002](https://doi.org/10.1016/j.swevo.2016.09.002).
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, USA.
- Regis, R.G., 2014. Evolutionary programming for high-dimensional constrained expensive black-box optimization using radial basis functions. *IEEE Trans. Evol. Comput.* 18, 326–347. doi:[10.1109/TEVC.2013.2262111](https://doi.org/10.1109/TEVC.2013.2262111).
- Sacks, J., Welch, W.J., Mitchell, T.J., Wynn, H.P., 1989. Design and analysis of computer experiments. *Stat. Sci.* 4, 409–435.
- Schmidt, M.D., Lipson, H., 2008. Coevolution of fitness predictors. *IEEE Trans. Evol. Comput.* 12, 736–749. doi:[10.1109/TEVC.2008.919006](https://doi.org/10.1109/TEVC.2008.919006).
- Settles, B., 2010. *Active Learning Literature Survey*. Technical Report 1648. Department of Computer Science, University of Wisconsin-Madison. Madison, WI, USA.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., de Freitas, N., 2016. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* 104, 148–175. doi:[10.1109/JPROC.2015.2494218](https://doi.org/10.1109/JPROC.2015.2494218).
- Siegmund, F., Ng, A.C.H., Deb, K., 2013. A comparative study of dynamic resampling strategies for guided evolutionary multi-objective optimization, in: *Proc. IEEE Congr. Evol. Comput.*, IEEE Press, Piscataway, NJ, USA. pp. 1826–1835. doi:[10.1109/CEC.2013.6557782](https://doi.org/10.1109/CEC.2013.6557782).
- Wang, H., Jin, Y., Doherty, J., 2017. Committee-based active learning for surrogate-assisted particle swarm optimization of expensive problems. *IEEE Trans. Cybern.* 47, 2664–2677. doi:[10.1109/TCYB.2017.2710978](https://doi.org/10.1109/TCYB.2017.2710978).
- Wang, H., Jin, Y., Sun, C., Doherty, J., 2018. Offline data-driven evolutionary optimization using selective surrogate ensembles. *IEEE Trans. Evol. Comput.* doi:[10.1109/TEVC.2018.2834881](https://doi.org/10.1109/TEVC.2018.2834881). in press.
- Wilson, W.R., Hay, M.P., 2011. Targeting hypoxia in cancer therapy. *Nat. Rev. Canc.* 11, 393–410. doi:[10.1038/nrc3064](https://doi.org/10.1038/nrc3064).
- Yu, J.C., Suprayitno, 2017. Evolutionary reliable regional Kriging surrogate and soft outer array for robust engineering optimization. *IEEE Access* 5, 16520–16531. doi:[10.1109/ACCESS.2017.2720631](https://doi.org/10.1109/ACCESS.2017.2720631).
- Yun, Y., Yoon, M., Nakayama, H., 2009. Multi-objective optimization based on meta-modeling by using support vector regression. *Opt. Eng.* 10, 167–181. doi:[10.1007/s11081-008-9063-1](https://doi.org/10.1007/s11081-008-9063-1).

Algorithm 1: Surrogate-assisted GA with pre-selection and static sampling

```
1  $N = 6, P = 20, k = 10, M = 1000, \mathcal{X} = 0.8, \mu = 1/N, s = 0.1, T = 3$ 
2 Initialise population  $Pop = \{\vec{x}_1, \dots, \vec{x}_P\}$  //  $\vec{x}$  normalised [-1,1]
3 Evaluate  $Pop$  with the real objective function  $k$  times and add to archive  $\mathcal{A}$ 
4 while evaluation budget not exhausted do
    /* build surrogate model */
5    Fit regression model  $\mathcal{R}$  using  $\mathcal{A}$ 
    /* pre-select offspring */
6    Parent  $p_1 \leftarrow \text{TournamentSelection}(Pop, T)$ 
7    Parent  $p_2 \leftarrow \text{TournamentSelection}(Pop, T)$ 
8    for  $M$  number of offspring do
9        Offspring  $a \leftarrow p_1$ 
        /* crossover */
10       if  $\text{Random}(0,1) < \mathcal{X}$  then
11           Perform uniform crossover with  $a$  and  $p_2$ 
        /* mutation */
12       for each parameter  $x$  in  $a$  do
13           if  $\text{Random}(0,1) < \mu$  then
14                $x \leftarrow x + \text{Random}(-s, s)$ 
        /* evaluate offspring with surrogate model */
15        $a.\text{utility} \leftarrow \text{Rating}(\mathcal{R}, a, \mathcal{A})$ 
        /* select, evaluate, and add the most promising offspring */
16       Evaluate the best utility offspring with the real objective function  $k$  times
17       Add offspring to  $\mathcal{A}$ 
18        $r \leftarrow \text{NegativeTournamentSelection}(Pop, T)$ 
19       Replace  $r$  with offspring
```

Algorithm 2: Expected improvement rating function

```
1 Input: fitted regression model  $\mathcal{R}$ , candidate  $a$ , evaluated archive  $\mathcal{A}$ 
2 Output: model expected improvement of  $a$ 
3  $ei \leftarrow 0$ 
4  $mean, std \leftarrow \mathcal{R}.predict(a)$ 
5 if  $std \neq 0$  then
6    $imp \leftarrow BestFitness(\mathcal{A}) - mean$  // minimising
7    $z \leftarrow imp/std$ 
8   /*  $cdf()$  is the standard normal cumulative distribution function */
9   /*  $pdf()$  is the standard normal probability density function */
10   $ei \leftarrow imp \times cdf(z) + std \times pdf(z)$ 
11 return  $ei$ 
```




Figure 1: Fittest individuals on the PhysiCell anti-cancer biorobots simulator. GA (triangle), GP-assisted GA (circle), and MLP-assisted GA (square).

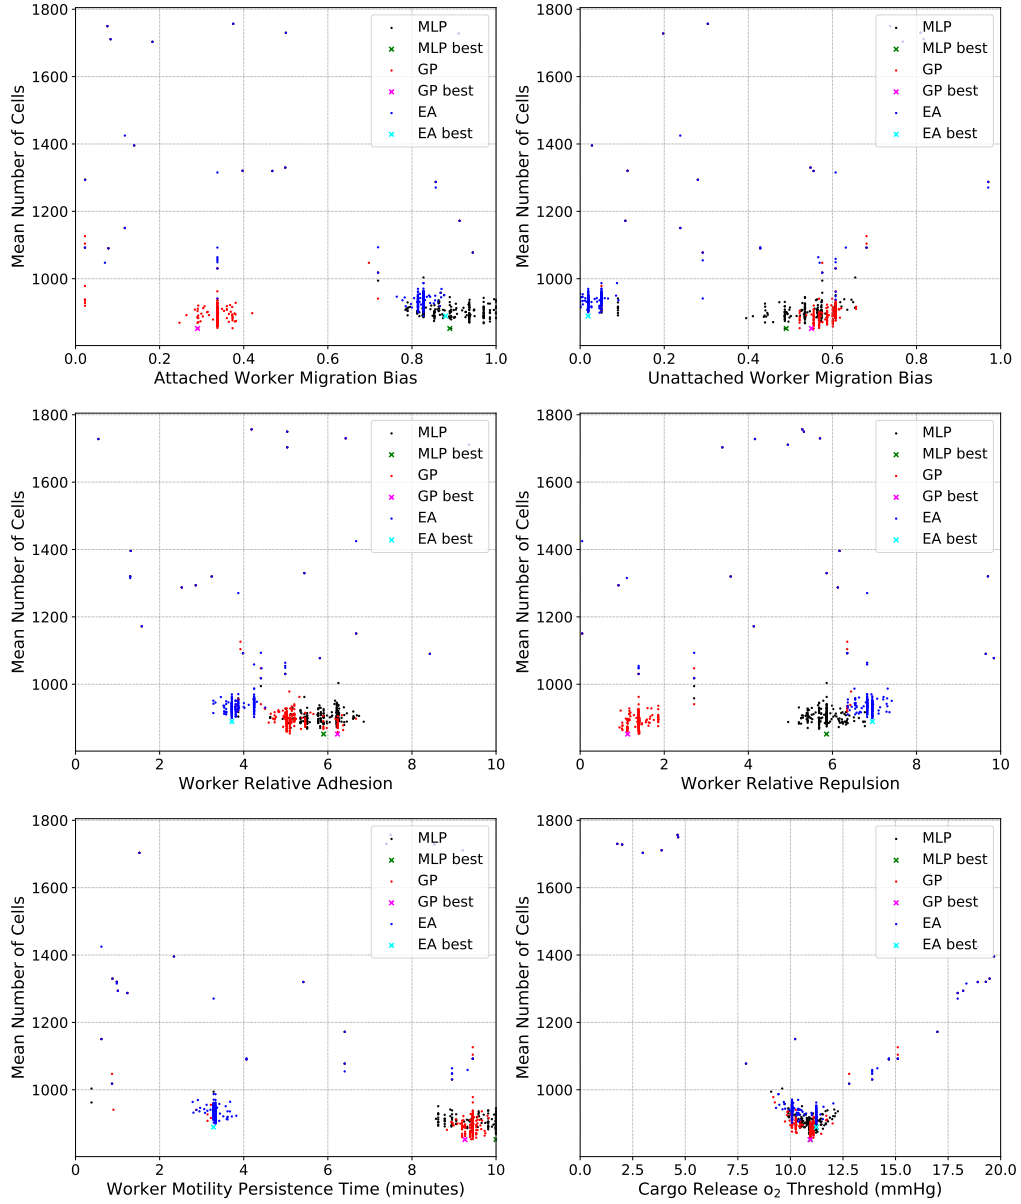


Figure 2: Scatter plot of all evaluated individuals on the PhysiCell anti-cancer biorobots simulator.

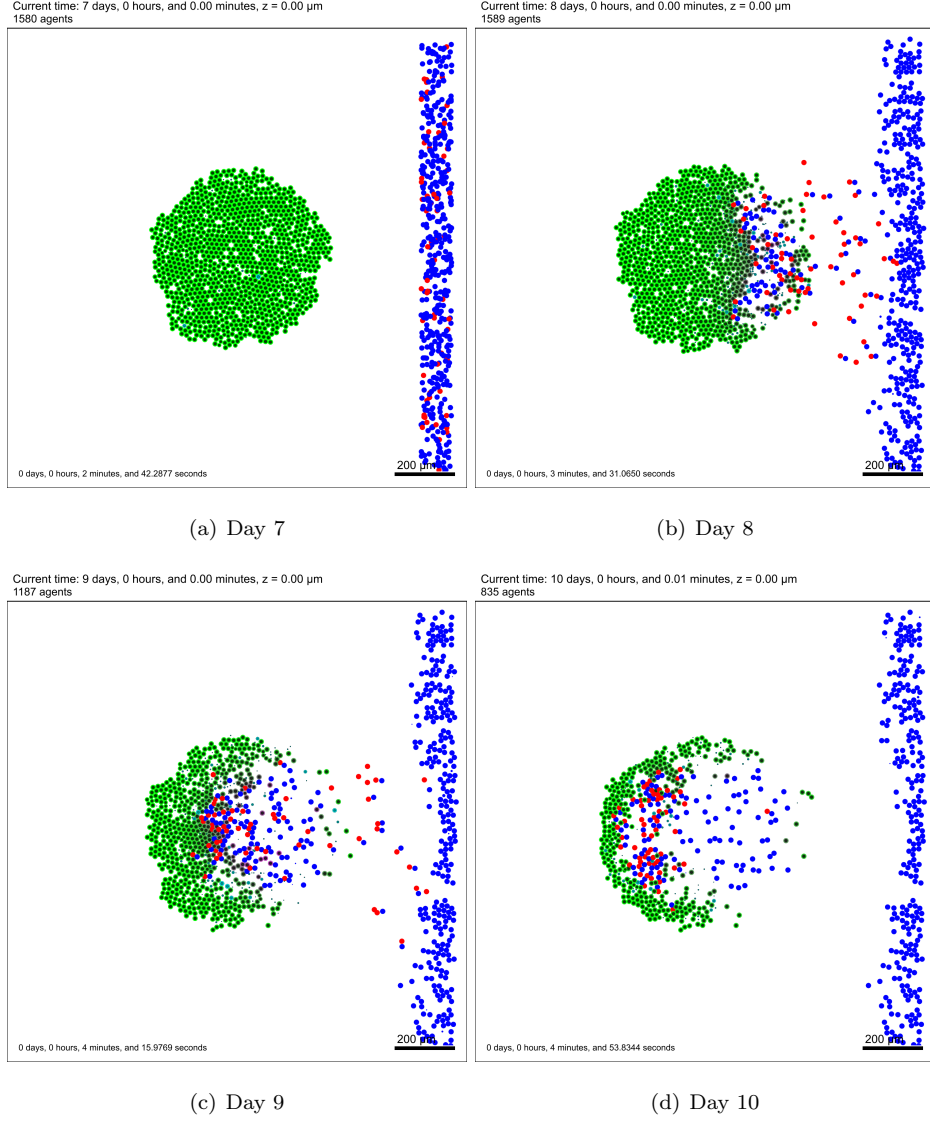


Figure 3: Example run of the fittest evolved individual with the GP-surrogate model on the PhysiCell anti-cancer biorobots simulator after 200 candidate evaluations (2000 simulations.) Shown are the worker cells (red), cargo cells (blue), and tumour cells (green) after 7 days of tumour growth and each subsequent day of treatment. Attached worker migration bias = 0.29; unattached worker migration bias = 0.55; worker relative adhesion = 6.24; worker relative repulsion = 1.13; worker motility persistence time = 9.26; cargo release o_2 threshold = 10.94. Mean number of cells after 10 simulated days = 852.