Cohort-Level Twinning in an Agent-Based Model of Metastatic Tumor Growth

Jan Schering[1 [0009−0009−6078−1965]], Lin Wouters[1 [0009−0004−0952−8378]], Inge Wortel[1 [0000−0003−3362−5229]], and Johannes Textor[1 [0000−0002−0459−9458]] ⋆

Data Science, Institute for Computing and Information Sciences, Radboud University, Nijmegen, The Netherlands

Abstract. Randomized controlled trials (RCTs) are a critical part of evidence-based medicine — they generate rigorous evidence supporting the efficacy of medical treatments at the population level. There has been much interest in performing RCTs in silico by creating a digital twin of each individual in a study population and using the twin to predict the effect of one or more candidate treatments of interest. However, this apparently straightforward approach relies on high-fidelity digital twins that are not usually available; indeed, thorough validation of digital human twins in interventional settings already requires RCT data. Here, we discuss a different approach: in cohort-level twinning, our goal is to create a digital representation of an entire patient cohort's baseline and response profiles, but not necessarily of each patient in the cohort itself. Cohort-level twinning allows us to leverage coarse-grained models such as ordinary differential equations and agent-based simulations to perform in silico trials that can be used to plan RCTs and make predictions about the population-level effects of novel interventions. We illustrate cohort-level twinning using an agent-based simulation of a metastasizing tumor implemented using a cellular Potts model. We fit this simulation to reproduce progression-free survival times in a cohort of non-small-cell lung cancer patients, and demonstrate how such a fitted model can be used to conduct an in silico trial and a power analysis.

Keywords: In silico trials · Approximate Bayesian Computation · Cellular Potts Model

1 Introduction

Cancer treatment has seen significant progress over the past decade, with an ever-increasing arsenal of different treatment modalities improving patient outcomes. Unfortunately, the rising cost and complexity of clinical trials have become a bottleneck in testing new (combinations of) treatments. This development has generated excitement for the prospect of conducting clinical trials *in silico*: clearly, it would be ideal if one could substitute costly, lengthy, and ethically sensitive *in vivo* interventions in human subjects with *in silico* interventions in computer

⋆ Corresponding author; e-mail: johannes.textor@ru.nl

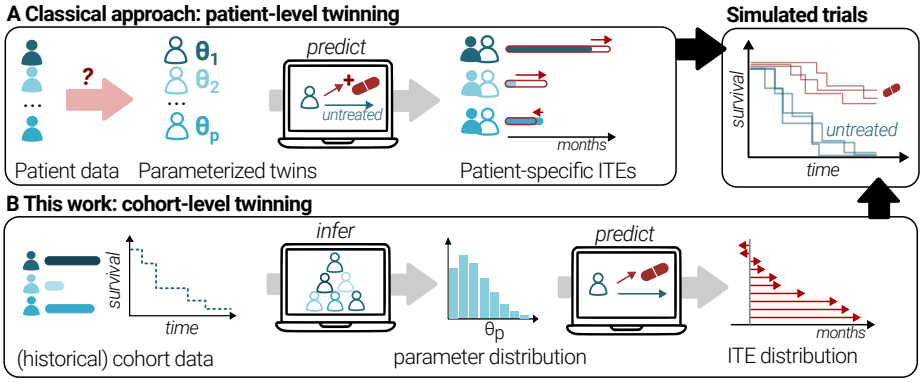

Fig. 1. Patient-level versus cohort-level twinning. A Most digital twin approaches focus on modelling individual patients. This requires the ability to derive model parameters θ_p from each patient's baseline data (red arrow), allowing the model to make patient-level predictions of clinical outcomes with and without treatment to derive "individual treatment effects" (ITEs) and simulate a clinical trial *in silico*. This approach has the advantage of providing ITE predictions for specific patients, but fails when there is insufficient knowledge or data to infer a patient's θ_p. **B** In cohort-level twinning, we instead infer a *distribution* of patient parameters to match historical cohort data, circumventing the need for directly predicting θ_p for individual patients. Although this does remove the ability to match ITEs to specific patients, a distribution of ITEs can still be modelled and used for *in silico* trials.

simulations (Fig. 1a). But perhaps the most significant advantage from an evidential perspective is a little more subtle: having a *digital twin* of each patient would make it possible to simulate both the baseline outcome and several interventions to get estimates of the personalized, *individual treatment effect* (ITE) for each patient-treatment combination. This approach seems to solve a fundamental problem of causal inference: we can never observe both outcomes – treated and untreated – for the same human at the same time. Therefore, randomized controlled trials can only estimate *average treatment effects* (expected benefit for the population at large) or *conditional average treatment effects* (expected benefit for different patient groups defined by covariates such as age).

Hence, digital twinning claims to solve an extremely challenging problem, but this comes at the cost of making very strong assumptions [3] – such as the ability to construct the digital twin in the first place. To get valid ITEs, digital twins must accurately predict *all potential outcomes*, including no treatment and all treatment options of interest, *for all individuals in the target population.* Thus far, we are not aware of any human digital twin that clears this bar in any disease setting. Even much simpler digital twins of single cells do not exist yet, despite decades of efforts to create them.

Recent work has shown that we may not always need a faithful digital twin to perform cancer therapy trials *in silico* [4, 6]. Instead, we can create what we call *cohort-level twins* (Fig. 1b): populations of simulation models with different

parameters that collectively produce a realistic outcome distribution for at least one arm of interest (e.g., for an untreated baseline population). A cohort-level twin will typically have the same size and possibly covariate profile as the target cohort of interest, but we cannot match each patient in the cohort twin to a patient in the real cohort, nor do we claim that predictions are necessarily accurate on the level of single simulated patients. Nevertheless, cohort-level twins have been shown capable of predicting important cohort-level outcomes – in technical terms, *distributions of* ITEs rather than individual ITEs – such as the "delayed curve separation" phenomenon commonly seen in cancer immunotherapy and the more immediate responses commonly seen in chemotherapy [4, 6]. While sacrificing the ability to predict personalized treatment effects for individual patients, this approach provides a useful framework for comparing the potential impact of hypothetical treatments targeting different biological mechanisms and for informing clinical trial designs.

Thus, by solving a less challenging – but still very important – problem, cohort-level twins require less restrictive assumptions than cohorts of patient-level twins, potentially making them more broadly applicable. However, this comes at the cost of a new challenge: cohort-level twins are parameterized by *fitting* parameter distributions to reproduce outcome distributions. This fitting approach is computationally intensive – it may, for example, require hundreds of thousands of simulations rather than just one simulation per patient. Especially for detailed simulations, this cost can be prohibitive.

In this paper, we show an example of cohort-level twinning for a moderately complex simulation of tumor growth and metastasis based on an agent-based model called the *cellular Potts model* [7, 11, 14]. We show how the model is constructed, how a caching strategy can ameliorate the computational cost of fitting, how to fit the model to a patient cohort, and how to perform a subsequent *in silico* trial for a hypothetical treatment.

2 Simulating Tumor Growth and Metastasis

Our model builds on the cellular Potts model (CPM) framework [7] to simulate a tumor growing stochastically from a single malignant cell. The CPM is spatial and describes cells c and extracellular matrix (ECM) as collections of voxels on a lattice (Fig. 2). CPMs simulate cell expansion, contraction, and emergent dynamics by sequentially attempting to copy the cell c occupying a random source voxel s into a neighboring target voxel t. These attempts are accepted or rejected based on their effect ΔH on a global energy H defined by the modeller.

As the sequential nature of this sampling algorithm would make the collection of a large simulated 3D dataset at high spatial resolution prohibitively slow, we here approximate it with a parallel Metropolis algorithm we implemented in TaichiLang [10], a Python platform that harnesses modern graphical processing units to enable large-scale physics simulations; the full source code for our model and all analyses in this paper is publicly available [12]. For each simulated Monte Carlo Step (MCS), we parallelize the copy attempts as follows (Fig. 2a):

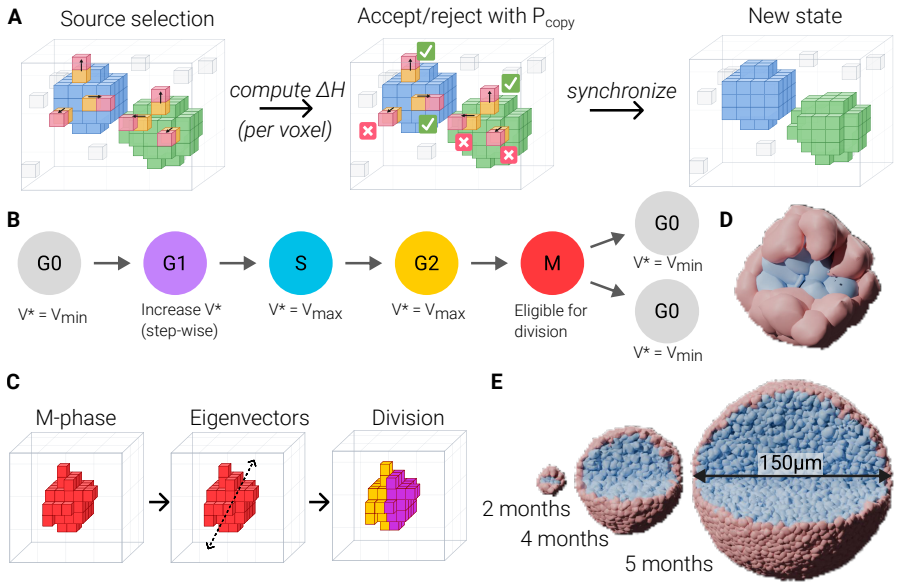

Fig. 2. A parallel CPM for tumor growth. A Approximate parallel Metropolis sampling of the copy attempts improves performance of the CPM. **B** Tumor growth is simulated through a cycle of cell growth and division. **C** Cells in M-phase are divided symmetrically using their eigenvectors. **D** Example tumor after 2 months with $g_0 =$ 10 days. Cells are smoothed for visualization only. Surface cells (red) are connected to the vascular system and can metastasize to other organs; internal cells (blue) cannot cause metastases in this model. **E** Tumor progression from the state shown in panel D.

1. Let every voxel i on the lattice be a simultaneous target t_i
2. Sample a neighboring source s_i for every target t_i
3. Calculate $\Delta H(s_i \to t_i)$
4. Accept or reject every proposed attempt with probability $P_{\text{copy}}(s_i \to t_i) = 1$ if $\Delta H \leq 0$, and $P_{\text{copy}}(s_i \to t_i) = e^{-\Delta H(s_i \to t_i)/\tau}$ otherwise.

The parameter τ adjusts overall stochasticity by scaling P_{copy}. $\Delta H(s \to t)$ is the change in system energy associated with the proposed copy attempt $s \to t$, obtained by comparing the lattice state before and after the update:

$$\begin{aligned} \Delta H(s \to t) &= \Delta H_{\text{volume}} + \Delta H_{\text{adhesion}} \\ &= \left(H^{\text{after}}_{\text{volume}} - H^{\text{before}}_{\text{volume}}\right) + \left(H^{\text{after}}_{\text{adhesion}} - H^{\text{before}}_{\text{adhesion}}\right), \end{aligned} \tag{1}$$

with H_{volume} the energy associated with the volume of the cell and H_{adhesion} the contact energy between cells, respectively defined as:

$$H_{\text{volume}} = \sum_{\text{cells } c} \lambda_V (V_c^{\text{current}} - V_c^*)^2 \; ; \; H_{\text{adhesion}} = \sum_{\text{voxel i}} \sum_{\text{neighbor j}} J_{ij}$$

Parameter	Value	Description
$V_{\min}, V_{\max}$	95, 190 voxels	Min. and max. per-cell target volume
λ_{V}	2	Scaling factor of the volume energy
Ω_{V}	155 voxels	Volume gate for cells advancing to S-phase
$J^{\mathrm{tumor}}_{\mathrm{tumor}}, J^{\mathrm{ECM}}_{\mathrm{tumor}}$	3, 6	Contact energies between tumor cells and ECM
g_0	[0, 30000] MCS	Mean length of G0, sampled per trial
g_1, S, g_2, M	94, 94, 37, 25 MCS	Mean length of each cycle phase
σ_{cycle}	0.1 MCS	Variance of the sampled phase lengths
N_x	29400 cells	Detection threshold
τ	10	Model temperature
time resolution	3 min/MCS	Conversion from MCS to minutes
spatial resolution	1 µm/voxel	Conversion from voxel to µm

Table 1. Parameters of the 3D Cellular Potts Model of tumor growth.

Here, V_c^* sets a cell c's target size, λ_V controls how much cells can deviate from that size, and J_{ij} is a contact energy between voxels (i, j): $J_{ij} = 0$ for voxels both belonging to the ECM or the same cell, $J_{ij} = J^{\mathrm{ECM}}_{\mathrm{tumor}}$ between tumor and ECM and $J_{ij} = J^{\mathrm{tumor}}_{\mathrm{tumor}}$ between two separate tumor cells.

Our proliferation model is built on an existing model of the cell cycle [11] in which cells pass through the 4 phases of cell division, with parameters g_1, S, g_2, and M describing the mean duration of each phase. This duration is stochastic and differs per cell, drawn from a uniform distribution bounded around its mean duration by a variance σ_{cycle}. We extend this model with a quiescent G0 phase with mean duration g_0 (Fig. 2b). Cells in G0 remain dormant around their natural volume $V_c^* = V_{\min}$. In the G1 phase, cells incrementally grow towards capacity $V_c^* = V_{\max}$ to prepare for mitosis. Progression to the S phase is gated by requiring a minimum volume Ω_{V}, simulating inhibition of growth when cells are compressed too much below their target volume. Cells stop growing in the S phase to simulate DNA synthesis, and once they progress through the G2 and M phases, divide perpendicularly to their minor axis (Fig. 2b,c). The resulting two daughter cells each start anew in G0. Most of these parameters can either be used from published work [11] or have little effect on tumor growth; for this proof of concept, we therefore focus on g_0 as the key parameter that controls outcomes in our simulations (Fig. 2d,e, Fig. 3a). All model parameters and descriptions can be found in Table 1.

Based on this tumor growth simulation, we model metastasis as follows. We start with a primary tumor $T1$ growing from a single cell with parameter g_0 (Fig. 3a). Once every simulated day, each surface tumor cell metastasizes with probability p_{met} to grow into a new metastatic tumor mass $T2$; here, we fixed $p_{\mathrm{met}} = 10^{-7}$ to capture the high overall probability of metastasis in our target cohort. Once $T1$ has reached the diagnosis threshold of 29400 cells (corresponding to a sphere with a diameter of 200 µm, a diagnosable microlesion [15]), it is detected and removed. Any metastasis $T2$ continues growing until it too reaches the detection threshold. Progression-free survival (PFS) is defined as the time between detection of $T1$ and $T2$ (Fig. 3b). For fitting purposes, we cap PFS at

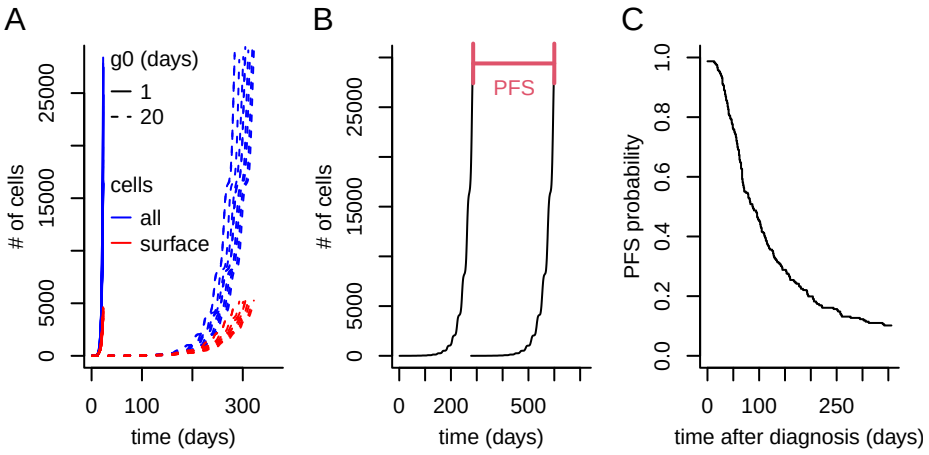

Fig. 3. Metastasis model. A Tumor growth curves for short and long g_0 simulated by the CPM; N=10 curves per g_0 value. **B** Primary tumor, metastasis, and definition of progression-free survival (PFS). **C** PFS times simulated by sampling g_0 from a lognormal distribution with median $e^{\mu} = e^{8.5}$ and multiplicative spread $e^{\sigma} = e^{1.15}$.

1 year, censoring metastasis detection at later times. To generate KM curves *in silico* based on this simulation, we draw parameter values from a suitable distribution for each simulated individual (Fig. 3c).

3 Fitting Stochastic Models to Population Data

To illustrate cohort-level twinning, we use the computational model to create a digital twin of the PROSE lung cancer patient cohort [8] (N=263, survival data from R package "kmdata" [9]). Since the CPM is still computationally expensive, and has to be run many times in the fitting procedure described below, we implemented a caching strategy to speed up simulations: a large "archive" of tumor growth curves for random g_0 phase durations between 0 and 62.5 days is pre-computed and cached. When conducting simulations at a given g_0, we sample growth curves for $T1$ and $T2$ from this cached archive instead of running the CPM (Fig. 4a).

Instead of measuring g_0 for individual patients, as we would have to do for patient-level twinning, we instead fit its *distribution* by finding a median e^{μ} and multiplicative standard deviation e^{σ} that produce samples ($g_0 \sim \exp(\mathcal{N}(\mu, \sigma))$) capturing the between-patient heterogeneity in PFS (Fig. 4b). We infer the parameters $\theta = (\mu, \sigma)$ of this distribution using *Approximate Bayesian Computation with sequential Monte Carlo* (ABC-SMC [13], Fig. 4c). Initially sampling from a uniform prior, our algorithm iteratively produces populations of N_{pop} accepted parameters θ that produce simulated outputs "close enough" to the data. As we are interested in matching progression-free survival (PFS), we here con-

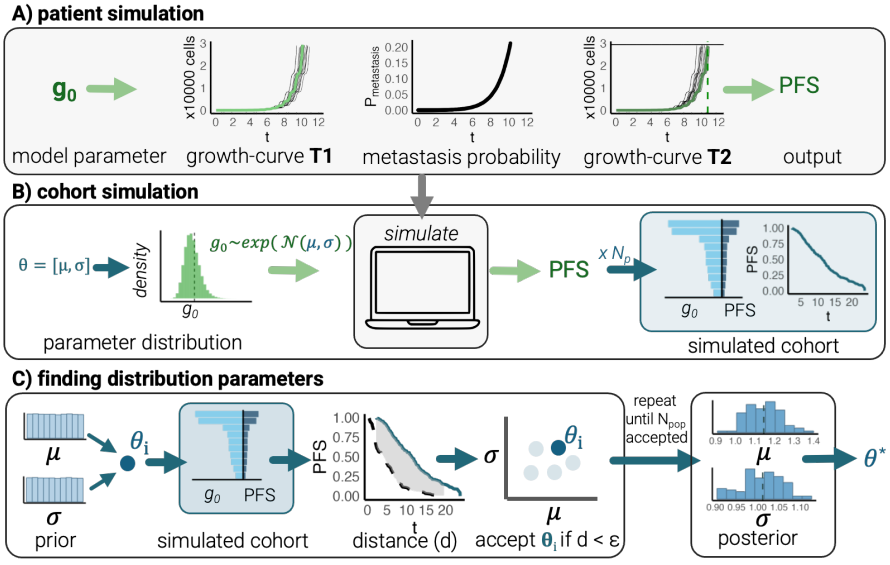

Fig. 4. Fitting the tumor simulator to cohort data. A To simulate PFS for parameter g_0, we sample a similar g_0^A from the cached archive to obtain the primary tumor growth curve ($T1$). With probability $P_{\text{metastasis}} = p_{\text{met}} \times N_{\text{surface cells}}$, a second growth curve for the metastasis ($T2$) is also sampled from the archive. $T2$ is detected when it reaches 29400 cells, yielding the PFS time. **B** A simulated *cohort* is created by sampling N_p values of g_0 from a lognormal distribution ($g_0 \sim \exp(\mathcal{N}(\mu, \sigma))$) and simulating the PFS for each sample as in A. **C** We fit the log-scale median μ and SD σ of the lognormal distribution over g_0 values using ABC-SMC. For each candidate pair $\theta_i = [\mu_i, \sigma_i]$ we create a simulated cohort, and accept the parameters if the distance between the experimental and simulated Kaplan-Meier curves is below the threshold ϵ_{round}. Once N_{pop} parameters are accepted, they become the prior for a next population; when sampling from these new accepted parameter values in the next round, a Gaussian noise term is added. Per round we use a lower ϵ_{round}. The population of accepted θ in the final round form a posterior sample, from which we can calculate the mode to get the final $\theta^* = (\mu^*, \sigma^*)$ at which to simulate the fitted Kaplan-Meier curves.

sider a simulated cohort an acceptable representation of the data when the root mean squared difference (RMSD) between Kaplan-Meier (KM) curves of patient and simulated cohorts is below threshold ϵ. By reducing ϵ for each subsequent ABC round, we progressively bring simulations closer to the data. The accepted θ in the final round approximate the posterior of θ given the data.

Fig. 5 illustrates the progressive improvement in simulated KM curves over ABC rounds in an example run. Sharp posteriors were found for both μ and σ, and over repeated runs, we achieved RMSDs of 0.04 or lower on the resulting KM curves. These results demonstrate that we can infer a distribution of g_0 to obtain a cohort-level twin that matches the PFS of real patient cohorts.

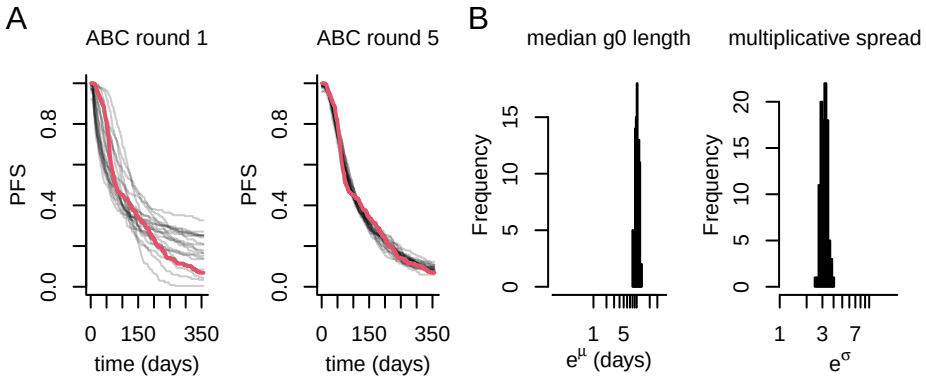

Fig. 5. Example ABC run fitting metastasis model to PROSE data. A KM curves corresponding to accepted θ in ABC rounds 1,2, and 5 (out of 5); **B** posterior distributions for μ and σ. In this run, we used uniform priors for θ over the X-axis ranges in the posterior plots ($\mu \sim U(4, 10)$ and $\sigma \sim U(0.1, 3)$). In each population, $N_{\text{pop}} = 100$ parameters were accepted. Following the initial population, candidate parameters were sampled from the previous population and perturbed by a Gaussian kernel with standard deviations $\sigma_\mu = 0.2$ and $\sigma_\sigma = 0.1$. We started with $\epsilon_1 = 0.3$ and reduced it every round to either half its value or the 10% quantile of the RMSDs from the accepted θ, whichever is higher. The algorithm was terminated when six populations (including the initial prior sample) were completed or when the algorithm was not making progress, defined as two consecutive populations rejecting more than 90% of the proposed parameter sets.

4 Predicting Treatment Effects in an In Silico Trial

Rather than replacing real clinical trials, an important application of *in silico* trials lies in aiding the planning and preparation of actual trials (we have used similar models to plan the MIND-DC clinical trial [2] and two additional trials due to start at Radboudumc Nijmegen at the end of 2026). The key idea is that *in silico* trials can help explore how a hypothetical treatment targeting a specific biological mechanism might translate into clinical outcomes. For example, how strong would the mechanistic effect need to be to provide clinically meaningful benefits? Would the resulting treatment effect violate common statistical assumptions such as proportional hazards? Which clinical endpoints would be most suitable, and how many patients would we need to achieve sufficient power?

To illustrate how our fitted simulation model can be used in this manner, we consider the fitted cohort as clinical baseline and model the effect of a treatment (e.g., an enhanced chemotherapy) that prolongs the G0 phase by a given factor when the primary tumor is diagnosed and treatment starts – e.g., for effect size 2, a tumor with $g_0 = 100$ would be set to $g_0 = 200$. Similar to previous work conducted using much simpler models [4], this allows us to predict the counterfactual KM curves for this hypothetical treatment (Fig. 6a). Here, our model predicts that prolonging the G0 phase by a factor 1.5-2 would yield noticeable

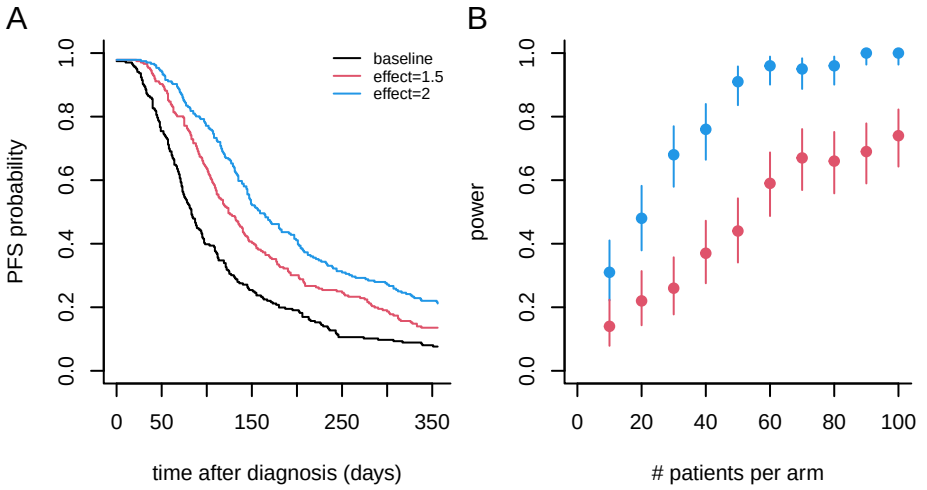

Fig. 6. In silico trial. A Predicting the effect of a treatment that prolongs g_0 by a given factor (effect) upon diagnosis of the primary tumor. **B** Power analysis using a logrank test comparing simulated cohorts with and without treatments (N=100 runs).

clinical effects with no clear violation of the proportional hazards assumption, suggesting that classical approaches such as a logrank test on PFS would be suitable to detect it. Fig. 6b shows a power analysis for our hypothetical treatment, indicating that 100 patients per arm would suffice to achieve 80% power for a therapy with an effect size of 2, but not 1.5.

5 Discussion

This paper provides a worked example of cohort-level digital twinning using a moderately complex simulation model fitted to survival data using ABC. We use the fitted model to predict treatment results for a hypothetical chemotherapy intervention and perform a power analysis as an example application.

The model used in this paper was intentionally limited given the available space and our intent to produce a clear example, and should be seen as a proof of concept only. Using such a model in practice would require further validation; for example, it would be good to demonstrate that the same model would still be able to fit a second external cohort with similar patient characteristics or a second clinical endpoint. Beyond validating the model on held-out data, it would be even more important to test sensitivity of model predictions to alternative choices in model structure and assumptions. For example, our model contains a fairly simple mechanism of cell growth and metastasis – but would alternative plausible models of this process yield similar predictions on the overall shape of PFS curves, suitable end points, and power analyses? What if we vary the overall length of the cell cycle, rather than only the G0 duration? For most applications, parts of the model structure and/or some of the parameters are likely

not completely identifiable based on available data (if they were, construction of patient-level twins would not be so challenging). Ideally, applications of cohort-level twins are therefore not based on any single model, but on several models, testing robustness of their predictions to account for this uncertainty. Indeed, even if models disagree on exact survival time predictions, they can still agree on their predictions on the level of trial design [4]. This is precisely what makes cohort-level twinning an attractive alternative to patient-level twinning: even if there remains too much uncertainty to make accurate patient-level predictions, some cohort-level predictions may be sufficiently robust to meaningfully inform trial designs.

Beyond varying the mechanisms of tumor growth and metastasis for the purpose of sensitivity analysis, a similar approach can be applied to more complex models. For example, it would be interesting to extend this simulation with a representation of the immune system, a key component of many existing immunotherapy models [5, 1], and aim to generate realistic data from the simulations for performing more advanced power analyses and other applications. By showcasing the general approach in a simple worked example, we hope to facilitate such efforts in future work.

Acknowledgments. The authors used Claude (Opus 4.8, Anthropic) and ChatGPT (GPT-5, OpenAI) to assist in drafting some of the code and figures. The authors reviewed and edited any generated content as needed and take full responsibility for the content of the publication.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bekker, R.A., Zahid, M.U., Binning, J.M., Spring, B.Q., Hwu, P., Pilon-Thomas, S., Enderling, H.: Rethinking the immunotherapy numbers game. Journal for ImmunoTherapy of Cancer **10**(7), e005107 (2022). https://doi.org/10.1136/jitc-2022-005107
2. Bol, K.F., Schreibelt, G., Bloemendal, M., van Willigen, W.W., Hins-de Bree, S., de Goede, A.L., de Boer, A.J., Bos, K.J.H., Duiveman-de Boer, T., Olde Nordkamp, M.A.M., van Oorschot, T.G.M., Popelier, C.J., Pots, J.M., Scharenborg, N.M., van de Rakt, M.W.M.M., de Ruiter, V., van Meeteren, W.S., van Rossum, M.M., Croockewit, S.J., Koeneman, B.J., Creemers, J.H.A., Wortel, I.M.N., Angerer, C., Brüning, M., Petry, K., Dzionek, A., van der Veldt, A.A., van Grünhagen, D.J., Werner, J.E.M., Bonenkamp, J.J., Haanen, J.B.A.G., Boers-Sonderen, M.J., Koornstra, R.H.T., Boomsma, M.F., Aarntzen, E.H.J., Gotthardt, M., Nagarajah, J., de Witte, T.J.M., Figdor, C.G., de Wilt, J.H.W., Textor, J., de Groot, J.W.B., Gerritsen, W.R., de Vries, I.J.M.: Adjuvant dendritic cell therapy in stage IIIB/C melanoma: the MIND-DC randomized phase III trial. Nature Communications **15**(1) (Feb 2024). https://doi.org/10.1038/s41467-024-45358-0
3. Cartwright, N.: No Causes in, No Causes out, p. 39–90. Oxford University PressOxford (Apr 1994). https://doi.org/10.1093/0198235070.003.0003
4. Creemers, J.H.A., Ankan, A., Roes, K.C.B., Schröder, G., Mehra, N., Figdor, C.G., de Vries, I.J.M., Textor, J.: In silico cancer immunotherapy trials uncover the consequences of therapy-specific response patterns for clinical trial design and outcome. Nature Communications **14**(1) (2023). https://doi.org/10.1038/s41467-023-37933-8
5. Creemers, J.H.A., Lesterhuis, W.J., Mehra, N., Gerritsen, W.R., Figdor, C.G., de Vries, I.J.M., Textor, J.: A tipping point in cancer-immune dynamics leads to divergent immunotherapy responses and hampers biomarker discovery. Journal for ImmunoTherapy of Cancer **9**(5), e002032 (2021). https://doi.org/10.1136/jitc-2020-002032
6. Creemers, J.H., Textor, J.: Leveraging mathematical models to improve the statistical robustness of cancer immunotherapy trials. Current Opinion in Systems Biology **40**, 100540 (Mar 2025). https://doi.org/10.1016/j.coisb.2024.100540
7. Glazier, J.A., Graner, F.: Simulation of the differential adhesion driven rearrangement of biological cells. Physical Review E **47**(3), 2128–2154 (Mar 1993). https://doi.org/10.1103/physreve.47.2128
8. Gregorc, V., Novello, S., Lazzari, C., Barni, S., Aieta, M., Mencoboni, M., Grossi, F., Pas, T.D., de Marinis, F., Bearz, A., Floriani, I., Torri, V., Bulotta, A., Cattaneo, A., Grigorieva, J., Tsypin, M., Roder, J., Doglioni, C., Levra, M.G., Petrelli, F., Foti, S., Viganò, M., Bachi, A., Roder, H.: Predictive value of a proteomic signature in patients with non-small-cell lung cancer treated with second-line erlotinib or chemotherapy (prose): a biomarker-stratified, randomised phase 3 trial. The Lancet Oncology **15**(7), 713–721 (2014). https://doi.org/10.1016/s1470-2045(14)70162-7
9. Guyot, P., Ades, A., Ouwens, M.J., Welton, N.J.: Enhanced secondary analysis of survival data: reconstructing the data from published kaplan-meier survival curves. BMC Medical Research Methodology **12**(1) (Feb 2012). https://doi.org/10.1186/1471-2288-12-9
10. Hu, Y., Li, T.M., Anderson, L., Ragan-Kelley, J., Durand, F.: Taichi: a language for high-performance computation on spatially sparse data structures. ACM Transactions on Graphics **38**(6), 1–16 (Nov 2019). https://doi.org/10.1145/3355089.3356506

11. Li, J.F., Lowengrub, J.: The effects of cell compressibility, motility and contact inhibition on the growth of tumor cell clusters using the cellular potts model. Journal of Theoretical Biology **343**, 79–91 (Feb 2014). https://doi.org/10.1016/j.jtbi.2013.10.008
12. Schering, J., Wouters, L., Wortel, I.M.N., Textor, J.: Cohort-level twinning in an agent-based model of metastatic tumor growth: simulation code (2026). https://doi.org/10.5281/zenodo.21702175
13. Sisson, S.A., Fan, Y., Tanaka, M.M.: Sequential Monte Carlo Without Likelihoods. Proceedings of the National Academy of Sciences **104**(6), 1760–1765 (2007). https://doi.org/10.1073/pnas.0607208104
14. Sultan, S., Devi, S., Mueller, S.N., Textor, J.: A parallelized cellular potts model that enables simulations at tissue scale (2023). https://doi.org/10.48550/ARXIV.2312.09317
15. Sun, J., Wu, S., Jin, Z., Ren, S., Cho, W.C., Zhu, C., Shen, J.: Lymph node micrometastasis in non–small cell lung cancer. Biomedicine & Pharmacotherapy **149**, 112817 (May 2022). https://doi.org/10.1016/j.biopha.2022.112817