

Frailty-Aware Dose Optimization in Geriatric Oncology: An MDP-Based Simulation Study

Panagiotis Karampelesis
otiskarampelesis@gmail.com
Dep. of Electrical and Computer
Engineering, University of Patras
Patras, Greece

Spyros Denazis
sdena@upatras.gr
Dep. of Electrical and Computer
Engineering, University of Patras
Patras, Greece

Evangelia I. Zacharaki
ezachar@upatras.gr
Dep. of Computer Engineering and
Informatics, University of Patras
Patras, Greece

Abstract

Older adults with cancer are underrepresented in computational oncology despite their increased vulnerability to treatment toxicity, frailty, and functional decline. In geriatric oncology, treatment planning must balance efficacy against quality of life, especially in patients with heterogeneous resilience and comorbidity profiles. In this paper, we present a Markov Decision Process (MDP)-based simulation framework for frailty-aware dose optimization in older breast cancer patients, grounded in the clinical protocol of the IMPORTANT project [4]. We generate synthetic longitudinal data including age, frailty, clinical test outcomes, dose level, and quality-of-life trajectories, and formulate a MDP to identify dosing policies that maximize cumulative quality-adjusted life years (QALYs). Simulation results suggest that adaptive, state-dependent treatment policies may outperform static full-dose or reduced-dose strategies, particularly in clinically vulnerable patients. These findings support the use of sequential decision models as a building block for digital twin-oriented decision support in geriatric oncology.

Keywords

Markov decision process, treatment optimization, Monte Carlo simulation, frailty, geriatric oncology

1 Introduction

Older adults with cancer represent a clinically vulnerable and highly heterogeneous population, yet they remain underrepresented in many computational oncology frameworks [6]. In geriatric oncology, treatment selection must balance efficacy against toxicity, frailty, functional decline, and quality of life. This challenge is particularly evident in older breast cancer patients receiving targeted therapy, where the optimal dose may depend not only on disease control but also on treatment tolerability and overall patient resilience. Moreover, older cancer patients often exhibit substantial variability in comorbidity burden, physiological reserve, and treatment tolerance. As a result, chronological age alone is insufficient for guiding treatment adaptation. Frailty-aware decision support is therefore of particular interest in geriatric oncology, where dose selection and subsequent modification may have a direct impact not only on tumor control but also on tolerability and long-term quality of life.

Sequential decision models offer a useful framework for this setting because they can explicitly represent evolving patient states and the long-term effect of repeated treatment choices. Such models are especially relevant when treatment adaptation depends on longitudinal information, including frailty and clinical status assessed over repeated follow-up visits. In this work, we present a

Markov Decision Process (MDP)-based simulation framework for frailty-aware dose optimization, informed by the decision logic of the IMPORTANT clinical study [4]. Using a synthetic longitudinal cohort of older breast cancer patients, we investigate how adaptive treatment policies can support individualized dosing decisions and improve quality-adjusted outcomes. More broadly, we discuss how this approach can serve as a computational step toward digital twin-oriented clinical decision support in geriatric oncology [5].

2 Methods

Our simulation setup is inspired by a clinical scenario in which older breast cancer patients undergo baseline assessment and comprehensive geriatric assessment (CGA) [7] before treatment selection. Patients are categorized as fit or frail, and treatment decisions are subsequently updated over repeated follow-up visits according to evolving clinical condition and vulnerability. The motivating clinical decision logic is based on a setting in which fit patients receive the recommended full dose, whereas frail patients are offered reduced-dose treatment and are monitored more carefully over time. During follow-up, treatment decisions may be maintained or adjusted according to updated clinical and geriatric information, reflecting the need for repeated reassessment in this vulnerable population [2].

We use Monte Carlo simulations to generate a synthetic longitudinal cohort of $N = 500$ patients followed for $T = 40$ quarterly visits. Ages were sampled between 71 and 95 years. Baseline frailty was initialized with higher prevalence for the frail group, consistent with the vulnerable population targeted by the motivating clinical scenario. For each visit, we simulated frailty status, clinical test outcome, treatment dose, and quality of life. The clinical test was modeled as a simple favorable/unfavorable indicator whose probability evolved according to previous status and treatment exposure. Frailty transitions were modeled as age-dependent and further modified by dose-related toxicity and clinical condition, allowing the cohort to capture plausible differences between fit and vulnerable patients over time. Quality of life was simulated conditional on frailty status, age, and treatment intensity, in order to reflect the trade-off between tolerability and therapeutic burden.

We then formulated a frailty-aware MDP in which each state encodes frailty status, recent clinical outcome, and dosing level [1, 3]. Available actions correspond to clinically relevant dose decisions (escalate/de-escalate), while rewards are defined in terms of QALYs (Figure 1). In the model development phase, data-driven probabilities are statistically approximated for transitioning between control and toxicity states given actions and patient features. This formulation allows the model to evaluate treatment benefit against toxicity

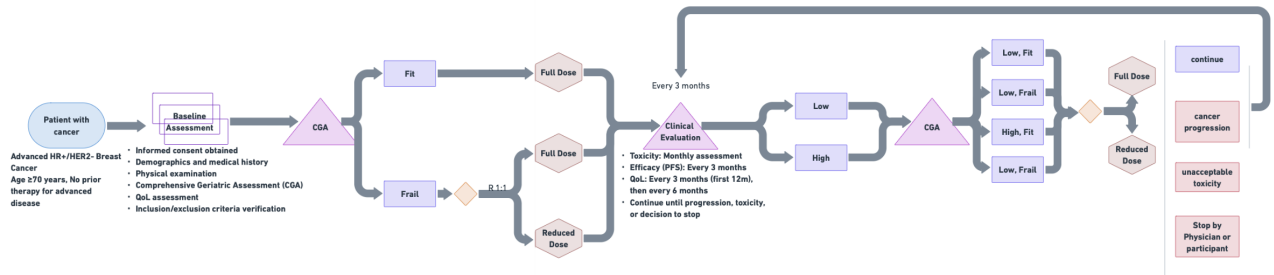


Figure 1: Clinical decision-support workflow aligned with the IMPORTANT study. Following baseline assessment and comprehensive geriatric assessment (CGA), patients are stratified into fit or frail groups, assigned an initial dosing strategy, and subsequently re-evaluated longitudinally using clinical and geriatric information to support adaptive treatment decisions.

burden and compute a policy that maximizes cumulative long-term outcome. In the prediction phase patient data are continuously integrated into the calibrated model to maintain an up-to-date digital twin (DT) of the patient. The DT enables real-time simulation of patient responses to alternative treatment strategies, with the inferred policy rules supporting treatment optimization.

3 Results

The simulated QALY trajectories of the different patient groups and treatment strategies are illustrated in Figure 2. The learned MDP policy on the synthetic data suggests different treatment behaviors across patient profiles. Fit patients tend to maintain their current dose until signs of frailty or reduced clinical tolerance emerge. In contrast, frail patients benefit from more adaptive moderation of treatment intensity, avoiding prolonged continuous exposure to full dose.




Figure 2: QALY trajectory analysis comparing the optimal MDP-derived dosing policy against fixed full-dose and reduced-dose strategies for frail and fit patients. The adaptive policy maintains more favorable QALY trajectories over time, with the clearest benefit observed in frail patients.

Compared with static full-dose and static reduced-dose strategies, the MDP-derived policy maintains more favorable per-cycle QALY trajectories. The effect is especially clear in frail patients, where the optimal policy separates early from the fixed strategies and preserves higher QALY values over time. In fit patients, the initial QALY level is higher, but the adaptive policy still mitigates the gradual decline associated with cumulative toxicity. Across the simulated cohort, these findings support state-aware dose individualization rather than one-size-fits-all treatment strategies.

More broadly, the simulation illustrates how geriatric assessment can be incorporated into a computational model to support treatment adaptation over time. This is particularly relevant in older adults, where the balance between efficacy and tolerability is highly patient-specific and may change throughout treatment.

4 Conclusion

We presented a synthetic data generation model and an MDP-based simulation study for frailty-aware dose optimization in older breast cancer patients. This work serves as a proof-of-concept, demonstrating how a Markov Decision Process can translate longitudinal patient assessments into adaptive treatment decisions. As this study relies on simulated environments, future work will focus on the integration of real-world clinical data and richer multimodal patient representations to formally assess the model’s clinical validity and impact.

Acknowledgments

The research conducted in this paper was supported by the European Union’s Horizon Research and420 Innovation Project, “Implementing geriatric assessment for dose Optimization of CDK 4/6-inhibitors in421 older bReasT cAncer patieNTs (IMPORTANT)” under Grant Agreement No. 101104589.

References

- [1] J. R. Beck and S. G. Pauker. 1983. The Markov process in medical prognosis. *Medical Decision Making* 3 (1983), 419–458.
- [2] E. J. Guerdard, A. M. Deal, Y. Chang, G. R. Williams, K. A. Nyrop, M. Pergolotti, et al. 2017. Frailty index developed from a cancer-specific geriatric assessment and the association with mortality among older adults with cancer. *Journal of the National Comprehensive Cancer Network* 15 (2017), 894–902.
- [3] F. Imani, Z. Qiu, and H. Yang. 2020. Markov decision process modeling for multi-stage optimization of intervention and treatment strategies in breast cancer. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 5394–5397.
- [4] IMPORTANT project. 2023. Implementing geriatric assessment for dose optimization of CDK 4/6-inhibitors in older breast cancer patients. <https://important-project.com/>. EU research project, duration 1 May 2023–30 April 2028.
- [5] E. Katsoulakis, Q. Wang, H. Wu, L. Shahriyari, R. Fletcher, J. Liu, et al. 2024. Digital twins for health: a scoping review. *npj Digital Medicine* 7 (2024), 77.
- [6] S. Pilleron and S. O’Hanlon. 2023. Digital twins for geriatric oncology: Double trouble or twice as nice? *Journal of Geriatric Oncology* 14 (2023), 101524.
- [7] T. J. Welsh, A. L. Gordon, and J. Gladman. 2014. Comprehensive geriatric assessment—a guide for the non-specialist. *International Journal of Clinical Practice* 68 (2014), 290–293.