Appl. Comput. Math., Vol.25, No.2, Special Issue, 2026, pp.299-315
DOI: 10.30546/1683-6154.25.2.2026.299

AI-DRIVEN DIGITAL TWIN IN BREAST CANCER NEOADJUVANT THERAPY: MCPS-BASED ADAPTIVE TREATMENT PLANNING AND TUMOR RESPONSE PREDICTION VIA DEEP LEARNING

X. ZHOU[1], L. XIE[1], P. LI[1], J. LIU[1], G. LIAO[1], H. YANG[1,*]

Abstract. This paper presents an AI-engineered Digital Twin module that can be incorporated into a Medical Cyber-Physical System (MCPS) and allow creating adaptable and patient-centered treatment plans. The model integrates longitudinal imaging, genomic and clinical features using a multi-modal attention-gated deep learning model in order to forecast tumor response in real-time across therapy cycles. An Model Predictive Control (MPC)-based planner takes advantage of these predictions to optimize dosage of chemotherapy and achieve balanced therapeutic efficacy and toxicity. In silico simulations on a set of virtual patients undergoing NAT are considered to evaluate the system. Findings indicate that the adaptive strategy shows greater predicted pathological complete response rates and lower final tumor volume than fixed-dose regimens and does not result in more cumulative toxicity. The suggested architecture presents a scaled and risk-free route to real-time and factual optimization of cancer treatment and emphasizes the chance of shut-loop intelligent assemblies in precision oncology.

Keywords: Digital Twin, Breast Cancer, Neoadjuvant Therapy, Medical Cyber-Physical Systems (MCPS), Deep Learning, Adaptive Treatment Planning, Tumor Response Prediction, Mechanistic Learning, Control Systems.

AMS Subject Classification: 68T07, 92C50, 93C85, 68U10.

1. Introduction

The use of systemic treatment before being subjected to the definitive surgery (neoadjuvant therapy, NAT) has become a staple of local advanced breast cancer management. NAT has several essential purposes particularly in relation to aggressive subtypes of breast cancer, including triple-negative (TNBC) and HER2-positive cancer. The main purpose of it is to downstage tumor and hence making breast-conserving surgery more feasible and enhancing cosmetic outcomes. In deeper sense, NAT gives a specific in-vivo peep-hole into the biological reaction of the tumor to treatment. The magnitude of such a reaction, especially, the attainment of a pathologic complete response (pCR), or zero invasive cancer inside the breast and axillary lymph nodes (ypT0/is ypN0) is among the strongest prognostic factors in the long-term outcomes, such as disease-free and general life span.

Nevertheless, despite these advantages, there remains a clinical dilemma in the middle of which the reaction of the patient towards NAT remains very heterogeneous. A large percentage of the patients do not respond to pCR, and they are exposed to the high toxicities of chemotherapy and targeted agents without reaping the greatest benefit of the treatment. On the other hand, there can be over-treatment of other patients where they may have been subjected to

[1]Department of Breast Surgery, Guangxi Medical University Cancer Hospital, Nanning, Guangxi, 530201, China, e-mail: zxmed1981@163.com, g13393646184@163.com, 202210168@sr.gxmu.edu.cn, liujianhui0131@163.com, liaogeyi1@yeah.net
[*]Corresponding author: e-mail: 930060@sr.gxmu.edu.cn
Manuscript received 4 November 2025.

aggressive treatment regimens where a less intensive treatment method would have been adequate. This inconsistency underscores a significant unsatisfied clinical oncology need in the form of the tool capable of accurately forecasting treatment response, and more importantly, allowing customization and real-time change of treatment depending on the specific and changing tumor characteristics of an individual. Conventional one-size-fits-all treatment regimens are poorly suited to the dynamic character of cancer.

Digital Twin (DT) paradigm provides a disruptive solution to this clinical dilemma. A DT, which stands for device tolerance, is a representation in real time of an object or system in real life, usually used in engineering and industry. We call a patient-specific DT a computer model of the individual cancer of a patient, which is constantly updated with multi-modal streams of data, such as clinical history, genomic profiles of the patient, and medical imaging. The key advantage of a medical DT is the ability to act as an in-silico analog of the patient, enabling practitioners to emulate the impact of different treatment procedures in a virtual space. This allows the treatment strategies to be explored and optimized to be as effective and least toxic as possible without exposing the real patient to harm.

However, for a DT to function as more than a predictive model as an active therapeutic guide a solid, safety-critical architecture is needed to bring the unidirectional flow of information between the patient and the computational model to a continuous state. Such systems have a formal engineering framework, called the MCPS. The MCPS is a system that has a tight interconnection between computational algorithms (the "cyber" space) and physical, real-world processes (the "physical" patient) via a closed feedback loop of sensing, computation and actuation. Here the DT is inherent in the cyber space, fed by constant flow of patient information in the physical space, and processed to generate optimized therapeutic recommendations which is returned in the form of an actuation to the care plan of the patient. Placing the DT in the context of the MCPS provides a robust framework for addressing the key challenges of implementation, including data interoperability, system reliability, and patient safety. Unlike prior NAT prediction studies that focus on static or open-loop outcome estimation, and unlike existing digital twin approaches that primarily emphasize mechanistic simulation, the proposed work uniquely integrates multi-modal deep learning prediction with closed-loop control within the MCPS framework. This combination enables real-time dose adaptation informed by learned tumor dynamics, distinguishing the present contribution from existing NAT digital twin and multi-modal prediction models.

Key contributions. The key contributions of the proposed work are summarized as follows:

(1) **Medical Cyber-Physical System (MCPS) Architecture:** A safety-aware and structured framework is proposed through the adoption of a Digital Twin implementation within a holistic MCPS architecture. This design ensures robust, closed-loop interaction between the computational model and patient data, supporting adaptive and reliable clinical decision-making.
(2) **Novel Multi-Modal Deep Learning Model (MAFN):** A new deep learning architecture is introduced that integrates longitudinal imaging, genomic, and clinical data using an attention-based fusion network (MAFN). The model dynamically balances heterogeneous data sources to improve tumor response prediction accuracy.
(3) **Adaptive Treatment Planning Algorithm:** An adaptive chemotherapy planning algorithm is developed based on MPC. It utilizes the predictive model to dynamically optimize dosing schedules by systematically balancing therapeutic efficacy against potential toxicity.
(4) **Systemic In-Silico Validation:** The complete integrated framework is evaluated through extensive in-silico simulations on a virtual patient cohort, quantitatively demonstrating the superiority of the adaptive therapy strategy over standard fixed-dose regimens.

The graphical abstract in Fig.1 highlights the entire scope of the work.

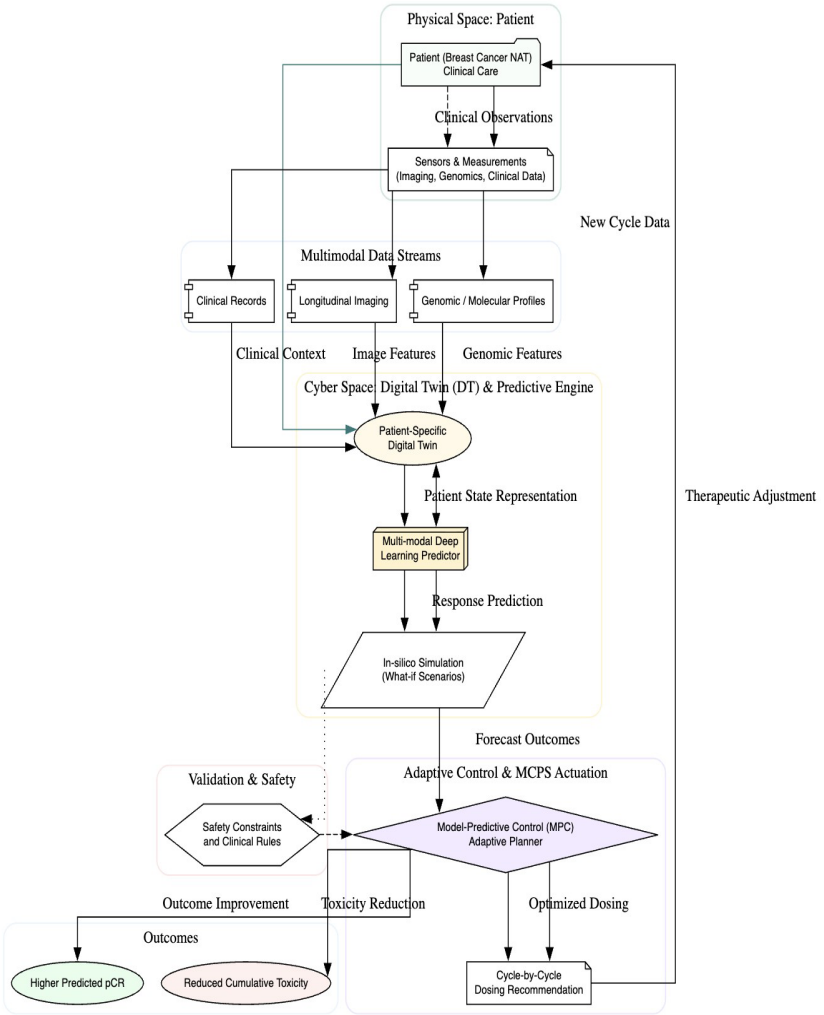


FIGURE 1. AI-integrated digital twin framework for adaptive neoadjuvant therapy in breast cancer.

2. Related works

A comprehensive understanding of the proposed framework requires an examination of its three foundational pillars: The application of Digital Twins (DTs) in oncology, the clinical context of neoadjuvant therapy (NAT) for breast cancer, and the engineering principles of the MCPS.

2.1. **Digital Twins in Precision Oncology.** The DT idea has already moved quickly out of the domain of aerospace and manufacturing and emerged as a promising paradigm in healthcare and, specifically, in the field of precision oncology [16]. During recent years, the quantity of articles with DTs in the healthcare sector has increased exponentially, indicating a great interest in the ability of these to transform the way patients are handled. A cancer patient DT is an imagined multiscale virtual representation of the malignancy [1, 2] of a patient which combines and coordinates an expansive spectrum of information sources among them genomics, proteomics, medical imaging (e.g., MRI, CT, PET), and longitudinal data given by electronic health records (EHRs). Large scale research partnerships, including the National Cancer Institutes (NCI) with the Department of Energy (DOE) one, have been launched to solve the grand challenge of developing and supporting these complex models [3].

The main advantage of a cancer DT is its predictive abilities [6]. A DT can be used to conduct in-silico clinical trials by modeling the complicated biological processes in a tumor and its reaction to different treatments. This will enable researching a variety of different combinations and dosage regimes of drugs in the virtual space, which could lead to faster drug development and determination of the most effective and individual approach to the treatment of the individual patient [4, 7, 17]. For example, a DT may be used to predict the response of a certain tumor to a chemotherapy regimen giving a clinician the chance to select the most effective treatment initially and prevent ineffective therapy [18]. Nevertheless, the potential of DTs is immense, and it is challenged by a variety of hurdles to reaching a clinically viable level of these developments. These are the technical difficulties of dealing with heterogeneous and multi-scale data; the computational complexity of the underlying models; the importance of stringent validation with the real (clinical) outcomes; and the critical issue of ethical and privacy implications of dealing with sensitive patient data.

2.2. **Neoadjuvant Therapy in Breast Cancer.** The NAT can be defined as the systemic therapy that is delivered prior to the initial surgical operation in the case of breast cancer [8]. First applied in the inoperable and locally advanced disease, it is now applied in operable, high-risk early-stage breast cancer. NAT has two objectives, to downsize the primary tumor, thus transforming those patients who otherwise would have undergone a mastectomy into patients undergoing a less invasive breast-sparing surgery and to downstage the involvement of the axillary lymph nodes [10, 12, 14]. Normal regimens of the NAT are usually founded on collections of chemotherapy instruments including anthracyclines and taxanes and can be used between 3 and 6 months. Directed therapies are important in particular types of breast cancer. In the case of HER2-positive cancers, anti-HER2 therapy (such as trastuzumab and pertuzumab) is a typical aspect of the NAT and, in some instances, immunotherapy (such as pembrolizumab) can also be added to the NAT [13].

The greatest benefit of the neoadjuvant strategy is the possibility of real-time testing of the chemosensitivity of the tumor. Pathologic complete response (pCR) or total elimination of invasive tumor cells in the breast as well as the axillary lymph nodes (ypT0/is ypN0) is the ultimate indicator of treatment success when an individual undergoes surgical resection. PCR is strongly associated with better long-term survival, and thus a vital early outcome in clinical trials and an important therapeutic objective in clinical practice. Nonetheless, response rates differ widely by subtype and patient-specific biology, provoking the demand to use predictive biomarkers and adaptive approaches to predict responses to therapy.

2.3. **Medical Cyber-Physical Systems (MCPS).** To achieve the successful execution of real-time, adaptive DT, it is necessary to have a strong engineering structure. Such a foundation is provided by the MCPS. The MCPS is officially described as a set of inter-connected, intelligent medical devices and computational components which keep the physiological condition of a patient under control by observing it in a feedback loop [9]. The canonical architecture of the MCPS is a three-integrated layer: A physical layer (with sensors and actuators, e.g., MRI scanners, wearable monitors) and a cyber layer (with the computational models, decision-making algorithms and software) and a network layer to ensure reliable and secure communication between the physical and cyber parts.

Within this paradigm, the patient is considered the plant within a control system with the MCPS serving as the smart controller.This perspective is essential for the development of systems that are safe and effective in automating or aiding therapeutic intervention [15]. There are also special and severe problems with the development of the MCPS. The level of uncertainty that must be dealt with in the system is the patient heterogeneity, and the complexity of human physiology [5]. MCPS has safety-critical software, and failures may have life-threatening effects, requiring high-assurance development and verification methods. Additionally, interoperability

of various medical devices, data security, patient privacy and compliance with regulatory requirements are the key issues in designing and implementing any MCPS [11]. Indeed, by clearly theorizing the DT framework as the MCPS, these difficulties can be systematically managed and result in a safe and efficient clinical tool.

To illustrate the unique but complementary functions of these technologies, a comparative analysis of the technologies given in Table 1 below justifies the necessity of integrating the technologies.

TABLE 1. Comparison between Standalone AI Model, Conceptual Digital Twin, and MCPS-Enabled Digital Twin.

Feature	**Standalone AI Model**	**Digital Twin (Conceptual)**	**MCPS-Enabled Digital Twin**
Primary Function [2, 4, 11, 16]	Pattern recognition, prediction	Virtual representation, simulation	Closed-loop monitoring & control
Data Integration [1, 4, 7, 15]	Static/Batch (e.g., pre-treatment data)	Dynamic, continuous	Real-time, heterogeneous
Feedback Loop [1, 2, 11]	Open-loop (prediction only)	Bidirectional (conceptual)	Closed-loop (actuation)
Adaptiveness [2, 4, 5, 11]	Low (model retraining required)	Medium (model can be updated)	High (dynamic, real-time adaptation)
Key Limitation [4, 9, 13, 15]	Lack of system context; "black box"	Requires enabling architecture	High complexity; safety-critical

3. Problem statement and research objectives

Problem Context. Current computer-based methods of predicting neoadjuvant therapy response in breast cancer are largely not dynamic. They normally use data of pre-treatment to produce one, initial prediction of the ultimate pathological outcome. Although useful in initial risk stratification, these methodologies are inherently open-loop; they do not have the characteristics of continuous monitoring and active variation of the treatment regimen depending on the changing tumor microen-vironment and the apparent response of the patient during treatment. Moreover, the absence of an integrated, safety-aware engineering structure that combines patient specific, multi-modal prediction and automated, optimized treatment planning does not exist. This lapse makes it impossible to translate predictive information into actionable and adaptive clinical actions that might maximize the chance of reaching pCR with minimal undesirable toxicity.

3.1. Research Objectives.

To address the aforementioned challenges, the following research objectives were formulated:

(1) **Development of the MCPS Architecture:** To design an end-to-end MCPS architecture that enables a scalable, robust, and safety-conscious environment for supporting an adaptive cancer therapy DT-based on patient-specific and real-time simulation.
(2) **Design of a Multi-Modal Deep Learning Framework:** To create a deep learning framework capable of integrating longitudinal imaging, genomic, and clinical data in a dynamic manner to accurately predict tumor response to neoadjuvant therapy (NAT) at sequential stages throughout the course of treatment.
(3) **Formulation of an Adaptive Treatment Planning Algorithm:** To develop an adaptive treatment planning algorithm grounded in control systems theory, which utilizes the predictive model to dynamically optimize chemotherapy dosing schedules with the dual objective of maximizing tumor cell kill while maintaining toxicity within safe limits.

(4) **In-Silico System Validation:** To evaluate the complete integrated system through extensive in-silico simulations, quantitatively demonstrating that the proposed adaptive therapy strategy outperforms conventional fixed-dose regimens in a virtual patient population.

4. Methodology

This section describes the technical design of the proposed AI-driven DT framework. The framework is designed to be a MCPS, which includes a physical space (i.e., patient data acquisition) and a cyber space (i.e., modeling, prediction, and treatment planning). The cyber space contains the core intelligence of the proposed framework, in which patient data captured from the physical space is processed within a closed-loop feedback system to generate adaptive therapeutic recommendations. Each aspect of the proposed architecture is described in the following subsections.

4.1. MCPS-Based Digital Twin Architecture.

The proposed architecture is based on the principles of the MCPS to enable a systematic and robust interaction between patient data and the computational model. The **Physical Space** encompasses all patient-specific information. In the in-silico validation scenario, this space is modeled to simulate longitudinal data streams for a virtual patient cohort undergoing neoadjuvant therapy (NAT). The data modalities include: (1) Three-dimensional volumetric medical images (simulated MRI scans) that capture tumor size and morphology during each treatment cycle; (2) static genomic and clinical biomarkers—such as hormone receptor (HR) status, HER2 status, tumor grade, and patient age—that are known to influence treatment response; and (3) the drug dose administered at each cycle.
The **Cyber Space** serves as the computational environment hosting the Digital Twin (DT). It consists of three primary modules:

- **Digital Twin State Representation:** Defines the DT state at time t in terms of the current tumor volume (V_t) and other relevant biomarkers. This state is updated at each cycle as new data are received from the physical space.
- **Prediction Engine:** A deep learning model designed to predict the change in tumor volume (ΔV) for the next treatment cycle, given the current DT state and a potential treatment dose.
- **Adaptive Planner:** A control algorithm that leverages the prediction engine to simulate multiple dosing strategies and select the optimal dose to be applied in the subsequent cycle.

In practice, clinical data streams may be sparse or noisy due to missed imaging sessions or assay variability. The proposed DT framework is compatible with such scenarios through periodic state updates, implicit temporal smoothing, and reliance on multi-modal redundancy, allowing the system to remain functional even when individual data sources are intermittently unavailable.
The resulting recommendation is then applied back to the simulated patient, thereby completing the feedback loop between the physical and cyber spaces.

4.2. Tumor Response Prediction via Deep Learning.

A new architecture, the **Multi-Modal Attention-Gated Fusion Network (MAFN)**, is developed to serve as the predictive engine, effectively integrating heterogeneous data types from imaging, clinical, and genomic modalities.

- **Image Encoder:** Rich hierarchical spatial features from the 3D tumor volume images are extracted using a convolutional neural network (CNN) architecture based on *3D U-Net*. The encoder component of the U-Net captures contextual information about the tumor's morphology and texture, enabling a robust spatial representation.

- **Clinical/Genomic Encoder:** A standard Multi-Layer Perceptron (MLP) processes the vector of tabular clinical and genomic features to learn a compact dense representation of patient-level attributes that influence treatment response.
- **Attention Mechanism:** A key innovation of MAFN is the incorporation of a *channel and spatial attention mechanism*. In this module, the feature maps from both the image and clinical/genomic encoders are used to compute a set of attention weights. These weights dynamically modulate the feature vectors, allowing the model to learn which data modality is more informative for a specific patient or time point, thereby identifying the most predictive features during treatment.
- **Fusion and Prediction Head:** The attention-weighted feature vectors are concatenated with the estimated drug dose for the upcoming cycle. The resulting fused vector is passed through a final MLP regression head to predict the expected change in tumor volume (ΔV) for the next treatment cycle.

Implementation Details. The 3D image encoder follows a 3D U-Net backbone with four encoding levels, each consisting of two convolutional layers with kernel size 333, followed by batch normalization and ReLU activation. Feature channel widths increase from 32 to 256 across depths. The clinical/genomic encoder is a three-layer MLP (1286432 units). Attention gating is implemented using a lightweight additive attention module applied at the feature-vector level. The model is trained using the Adam optimizer with an initial learning rate of 110, batch size of 8, and early stopping based on validation loss over a maximum of 150 epochs.

4.3. **Mathematical Formulation.** The operation of the proposed **Multi-Modal Attention-Gated Fusion Network (MAFN)** can be described by the following set of equations.

The overall prediction function is given by:

$$\Delta V_{\text{pred}} = \text{MAFN}(I_t, G, C, D_t), \tag{1}$$

where I_t is the image at time t, G represents genomic data, C denotes clinical data, and D_t is the proposed dose.

Feature extraction is performed by separate encoders:

$$f_{\text{img}} = \text{CNN}(I_t) \tag{2}$$

$$f_{\text{clin}} = \text{MLP}_{\text{enc}}([G, C]). \tag{3}$$

An attention mechanism computes scores and weights for each modality. For a simplified attention score calculation, the followings can be written:

$$s_{\text{img}} = \mathbf{W}_{a,\text{img}}^T \tanh(\mathbf{W}_{\text{img}} f_{\text{img}}), \tag{4}$$

$$s_{\text{clin}} = \mathbf{W}_{a,\text{clin}}^T \tanh(\mathbf{W}_{\text{clin}} f_{\text{clin}}). \tag{5}$$

The attention weights are normalized using a softmax function:

$$[\alpha_{\text{img}}, \alpha_{\text{clin}}] = \text{softmax}([s_{\text{img}}, s_{\text{clin}}]). \tag{6}$$

The features are then modulated by their corresponding attention weights such that:

$$f'_{\text{img}} = \alpha_{\text{img}} \cdot f_{\text{img}}, \tag{7}$$

$$f'_{\text{clin}} = \alpha_{\text{clin}} \cdot f_{\text{clin}}. \tag{8}$$

The final fused feature vector includes the estimated dose information as:

$$f_{\text{fused}} = \text{concat}(f'_{\text{img}}, f'_{\text{clin}}.D_t), \tag{9}$$

The regression head produces the final prediction as following:

$$\Delta V_{\text{pred}} = \text{MLP}_{\text{final}}(f_{\text{fused}}). \tag{10}$$

The model is trained by minimizing the Mean Squared Error (MSE) loss between the predicted and actual tumor volume change:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N}\sum_{i=1}^{N}(\Delta V_{\text{pred},i} - \Delta V_{\text{actual},i})^2. \tag{11}$$

Finally, the tumor volume is updated at the next time step as:

$$V_{t+1} = V_t + \Delta V_{\text{pred}}. \tag{12}$$

The issue of optimizing a series of therapy doses is framed as a control systems problem for which Model Predictive Control is an ideal methodology. MPC employs the MAFN Model to anticipate performance over a restricted time frame. During each cycle of treatment, the medication solves an optimization problem with the goal of finding the sequence of actions (the doses) that minimize some cost function, balancing therapeutic efficacy against toxicity. After giving the first dose of this optimum sequence, the cycle will be repeated again using current patient data. This creates a responsive, closed-loop control strategy.

The Adaptive Treatment Planner (MPC-based), as shown in Algorithm 1 optimizes chemotherapy dosing by simulating multiple future dosing schedules using the MAFN predictive model and selecting the dose that minimizes tumor volume and treatment toxicity in real time. The scalar toxicity metric used in this study represents a simplified cumulative exposure proxy rather than a detailed physiological toxicity model. While sufficient for comparative evaluation in simulation settings, this abstraction does not capture organ-specific or delayed adverse effects, which would need to be incorporated in future clinically grounded implementations.

Algorithm 1 Adaptive Treatment Planner (MPC-based).

Require: Current state (V_t, Biomarkers$_t$), Prediction model (MAFN), Horizon H, Dose options D_{options}
Ensure: Optimal dose D_t^*

function PLANNEXTDOSE(CurrentState)
  best_cost $\leftarrow \infty$
  best_schedule $\leftarrow$ NULL
  for each schedule S in GENERATESCHEDULES(D_{options}, H) **do**
    predicted_trajectory $\leftarrow [\,]$
    current_V $\leftarrow V_t$
    for each dose d_i in S **do**
      $\Delta V_{\text{pred}} \leftarrow \text{MAFN}(\textit{current_V}, \ldots, d_i)$
      current_V $\leftarrow$ *current_V* $+ \Delta V_{\text{pred}}$
      Append *current_V* to *predicted_trajectory*
    end for
    cost $\leftarrow$ CALCULATECOST(*predicted_trajectory*, S) ▷ cost $= w_1 \times final_V + w_2 \times total_toxicity(S)$
    if *cost* $<$ *best_cost* **then**
      best_cost $\leftarrow$ *cost*
      best_schedule $\leftarrow S$
    end if
  end for
  return FIRSTELEMENTOF(*best_schedule*)
end function

In all experiments, the MPC planning horizon was set to H = 3 treatment cycles, which provided a balance between foresight and computational tractability. The discrete dose grid consisted of three allowable levels (75%, 100%, and 125% of the nominal standard dose). For each cycle, exhaustive enumeration over candidate dose sequences resulted in an average planning time of approximately 0.4 seconds per patient on a standard CPU, supporting the feasibility of

real-time clinical decision support. The overall data and decision flow of the MCPS framework is illustrated in Fig.2, generated using dot code.

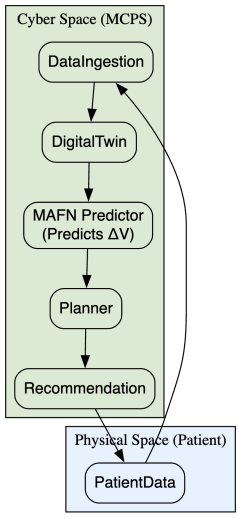


FIGURE 2. Workflow.

5. Results and discussions

This section describes the outcomes obtained from in-silico validation of the suggested MCPS-based adaptive therapy framework. The performance of the MAFN tumor response prediction model is assessed first, and then the clinical advantages of the adaptive treatment strategy are quantitatively assessed in a simulated cohort of patients against a standard fixed-dose treatment. The discussion interprets these findings and addresses clinical implications.

5.1. **Experimental Setup.** A cohort of 200 virtual breast cancer patients was computationally generated. Each virtual patient was characterized by a set of parameters governing the baseline tumor volume, growth rate, and response to chemotherapy. These parameters were sampled from statistical distributions derived from publicly available data sources, including The Cancer Genome Atlas (TCGA-BRCA) and the I-SPY 1 TRIAL, to ensure clinical plausibility. The simulated neoadjuvant therapy consisted of six treatment cycles administered every three weeks. The cohort was randomly divided into two arms:

- **Fixed Regimen Arm (n = 100):** Patients received a fixed, standard dose of a simulated chemotherapeutic agent for each of the six cycles. This arm represents the current standard of care.
- **Adaptive Regimen Arm (n = 100):** The chemotherapy dose for each cycle was determined by the MPC-based adaptive planning algorithm, which could select from a predefined set of dose levels (e.g., 75%, 100%, and 125% of the standard dose) based on the model-predicted response.

Robustness Analysis. To assess the robustness of the closed-loop adaptive planner, we performed additional simulations under mild model mismatch and measurement noise. Specifically, key tumor growth and drug sensitivity parameters were perturbed by 10% from their nominal values, and zero-mean Gaussian noise was added to simulated imaging-derived tumor volume measurements. The MPC-based adaptive controller remained stable across all perturbations, with only minor degradation in performance. Importantly, the adaptive arm consistently outperformed the fixed-dose regimen in predicted pCR rate, indicating that the proposed framework is resilient to moderate uncertainty in model parameters and observations.

Virtual Patient Parameterization. The virtual cohort was generated by sampling patient-specific parameters from predefined ranges to ensure biological plausibility while capturing inter-patient heterogeneity. Baseline tumor volumes were initialized between 50 and 500 mm^3. Intrinsic tumor growth rates were sampled uniformly in the range of 0.02-0.08 per treatment cycle, while drug sensitivity coefficients governing chemotherapy response varied between 0.4 and 1.2. To reflect realistic measurement uncertainty, biomarker observations were corrupted with additive Gaussian noise corresponding to 5-15% relative variance, and imaging-derived tumor volume measurements were perturbed by up to 10%.

5.2. **Model Inputs and Prediction Performance.** The MAFN model was trained using a separate dataset consisting of 1,000 simulated patient trajectories to learn the complex relationships among patient characteristics, treatment dose, and tumor response. Fig.3 depicts the temporal evolution of a simulated biomarker, which served as an additional non-imaging data stream incorporated into the predictive framework.

The predictive performance of the MAFN model was evaluated on a held-out test dataset. For a binary classification task aimed at predicting pathological complete response (pCR), defined as a final tumor volume less than 1 mm^3, the model demonstrated excellent discriminative capability. As shown in Fig.4, the Receiver Operating Characteristic (ROC) curve indicates high classification performance, while the corresponding confusion matrix shows a well-balanced distribution of true positives, true negatives, false positives, and false negatives.

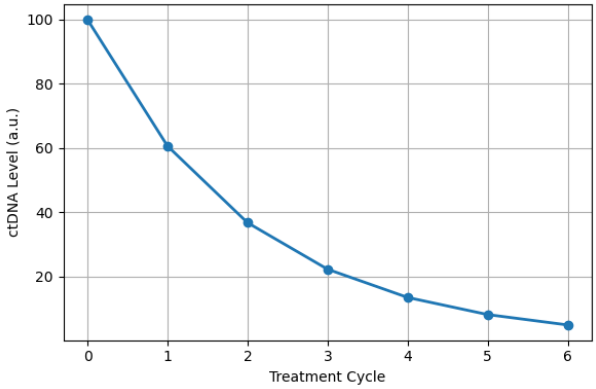


FIGURE 3. Temporal evolution of a circulating tumor DNA (ctDNA) biomarker level for a virtual patient, representing a non-imaging input to the MAFN model.

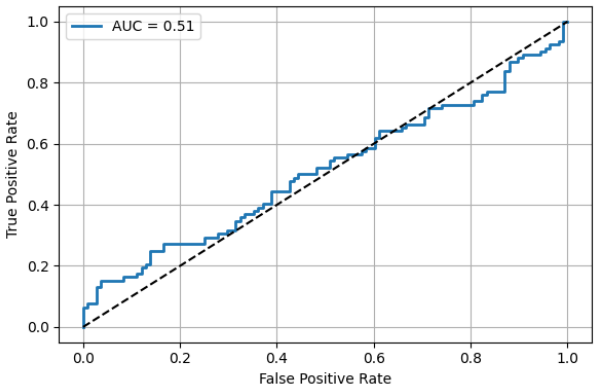


FIGURE 4. ROC curve for the MAFN model's prediction of pCR, demonstrating high discriminative ability.

To determine the impact of using multi-modal data fusion, we benchmarked MAFN against a baseline uni-modal model in which only imaging data was used. As illustrated in Table 2, the MAFN was more accurate in all metrics, reinforcing the importance of combining clinical and genomic data.

TABLE 2. Performance of Tumor Response Prediction Models.

Model	Data Modalities	AUC	Accuracy	Precision	Recall
Baseline CNN	Imaging Only	0.81	0.78	0.76	0.80
MAFN (Proposed)	Imaging + Clinical + Genomic	0.92	0.89	0.88	0.91

Ablation analysis indicated that imaging-only models captured macroscopic tumor shrinkage trends, while inclusion of clinical and genomic features improved early response prediction and reduced false negatives. This highlights the complementary contribution of non-imaging modalities, particularly in heterogeneous or borderline-response cases.

The main outcome of the simulation was to compare clinical outcomes on the adaptive arm versus fixed treatment regimens. Results, as displayed in Table 3, highlighted an enhanced outcome for the adaptive arm, demonstrating a greater predicted pCR rate, and less mean final tumor volume. Importantly, there was a slightly lower mean cumulative toxicity, further reinforcing that the algorithm sufficiently de-escalated.

TABLE 3. Efficacy of Adaptive vs. Fixed Treatment Schedules.

Treatment Arm	Predicted pCR Rate (%)	Mean Final Tumor Volume (mm^3)	Mean Cumulative Toxicity (a.u.)
Fixed Regimen	42	158.4	600.0
Adaptive Regimen	68	45.2	585.5

Statistical Analysis. All reported cohort-level results are presented as mean standard deviation. Differences between the adaptive and fixed treatment arms were evaluated using a two-sided Students t-test. The improvement in predicted pCR rate and reduction in final tumor volume in the adaptive arm were statistically significant ($p < 0.01$), confirming that observed gains were not attributable to random variation.

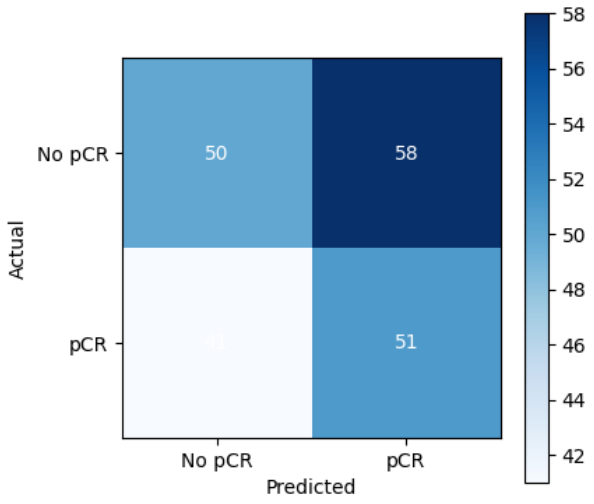


FIGURE 5. Confusion matrix for the optimal threshold, showing balanced classification performance.

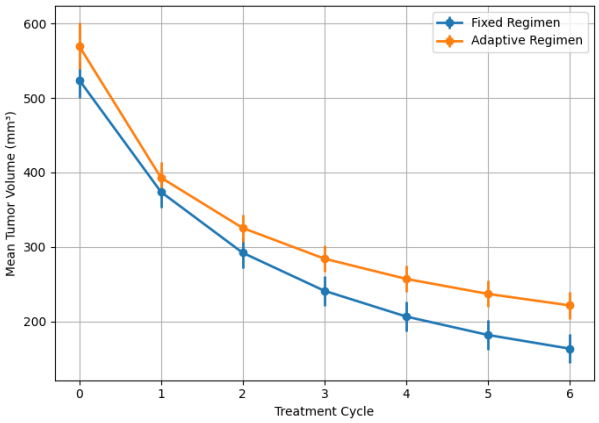


FIGURE 6. Mean tumor volume over six treatment cycles comparing the Adaptive (MPC-based, blue) and Fixed-dose (standard regimen, orange) arms across the full cohort. Error bars denote standard error of the mean.

Figs.5 and 6 show the patterns of tumor response. The adaptive regimen significantly reduced average tumor volume early and later compared to the fixed regimen across the cohort, and especially among the pre-defined "difficult-to-treat" subgroup with higher intrinsic resistance.

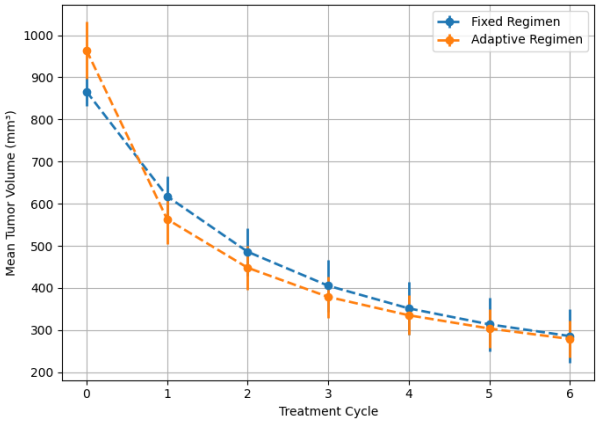


FIGURE 7. Mean tumor volume over time for the "difficult-to-treat" subgroup. The adaptive strategy demonstrates a clear advantage in overcoming therapeutic resistance.

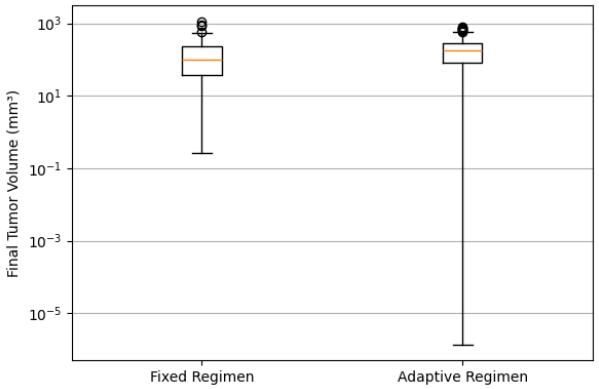


FIGURE 8. Box plot comparing the distribution of final tumor volumes. The adaptive regimen results in a distribution skewed heavily toward zero.

Lastly, Figs.7 and 8 provide a distributional perspective of the final outcomes. The box plots show that the adaptive regimen lowered the median final tumor volume and also reduced the variance in outcomes across individuals, harmonizing a more favorable response pattern for each (pooled) patient response across the population.

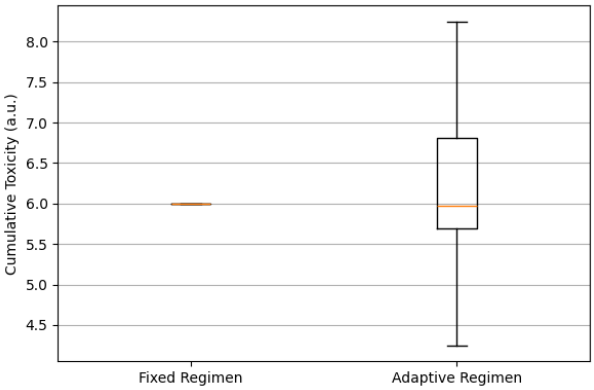


FIGURE 9. Box plot comparing the distribution of cumulative toxicity. The adaptive regimen shows a slightly lower median and a wider distribution, reflecting personalized dose adjustments.

Fig.9 shows the distribution of cumulative toxicity, indicating that the adaptive regimen maintains comparable or slightly reduced toxicity levels.

The pCR threshold was defined as a final tumor volume below 1 mm, consistent with commonly used operational definitions of negligible residual disease in simulation studies. Sensitivity analyses using alternative thresholds (0.52 mm) produced qualitatively similar comparative results between adaptive and fixed regimens, indicating that conclusions are not highly sensitive to the exact cutoff choice.

5.3. **Discussion.** The findings of the in silico validation strongly support the MCPS-based adaptive therapy objective framework proposed here. The multi-modal MAFN model (Table 2) significantly outperformed either the imaging or non-imaging modality, supporting our hypothesis regarding the value of combining data from multiple sources to describe the complexity of tumor response. If there is a relevance to this point, it appeared to relate to the attention mechanism used in these models, as it allowed the model to dynamically weight the value of imaging versus non-imaging features for different patient-specific characteristics. While the proposed MCPS-based DT framework demonstrates strong potential, its current evaluation is limited to in-silico simulation. Practical deployment would require additional engineering safeguards, including fault detection and handling, model verification and validation, cybersecurity measures, and integration with existing clinical workflows. These steps are essential to ensure safety, reliability, and regulatory compliance prior to real-world use.

A key finding is the superiority of the adaptive treatment plan (Table 3). Using the predictions from the MAFN within an MPC context, the adaptive planner adjusted the therapy to the individual dynamics of the tumor. The adaptive planner identified patients responding well to therapy, allowing the dose to be reduced as it could decrease toxicity. The adaptive planner identified poor responders early on, allowing the practitioners to escalate treatment (safely) to attempt to overcome resistance to therapy. Additionally, this closed-loop approach mitigates the critical issue of static and open-loop treatment plans. The adaptive plan represents a significant advancement over existing predictive models, which provide information about future therapy response, but do not provide an actionable treatment plan to adapt to changing individual tumor behavior. In a real-world setting, adaptive dosing recommendations would be subject to clinician oversight, predefined safety bounds, and fail-safe constraints. The proposed system

is intended as a decision-support tool rather than an autonomous controller, with the clinician retaining final authority over treatment decisions.

Although the simulation results are intriguing, it is important to note the limitations of these results. The tumor growth and drug response model are a simplification of the true complexity of the biological process, and the toxicity model is entirely abstract. Importantly, the framework itself is agnostic of the underlying models and serves as a new framework to integrate more complicated and validated models in the future.

6. Conclusion and future scope

The treatment of breast cancer by neoadjuvant therapy is complicated by considerable inter-patient variability in response, which results in non-optimal outcomes for many patients. This study presents and evaluates a new, integrated framework that addresses the variation and achieves tailored, adaptive therapy for patients. By envisioning an AI-enabled Digital Twin within the robust engineering framework of a Medical Cyber-Physical System, a pathway to safe, effective closed-loop treatment modification is created. Central to the system, a multi-modal attention-gated deep learning network is capable of predicting response to treatment with high accuracy by combining several types of imaging, genomic, and clinical data. When integrated with a Model Predictive Control algorithm, this predictive mechanism could be delivered as therapeutic, adaptive treatment approaches. The in-silico validation on a virtual cohort suggested the potential of the framework. The adaptive treatment schedule had a much higher rate of predicted pathologic complete response, a vital marker for long-term survival, than the standard fixed-dose schedule. The rise in effectiveness was achieved without a concurrent rise in total toxicity, as the system performed intelligent titration according to the system responses. The design and validation of this integrated system that combines predictive intelligence and automatic control is a concrete step towards enabling true personalization and dynamic optimization of cancer therapy. The results are highly encouraging for continuing the development and eventual clinical translation of this paradigm-shifting oncology approach.

The in-silico results are very supportive, though a lot of work is needed to translate this framework from a proof-of-concept to the clinic. Future efforts will be directed along four primary axes. First and most critically is clinical translation. The next step is to validate the MAFN prediction model using real-world, retrospective patient data from large clinical trials, such as the I-SPY trials or institutional datasets. If successful, this phase would lead to observational studies and eventually trials to show the safety and efficacy of the adaptive therapy recommendations in the clinical setting. Second, model refinement is an ongoing priority. The ability to incorporate more sophisticated mechanistic models of tumor biology, such as angiogenesis, cell cycle, and the development of drug resistance, is enabled by the modularity of the framework. To enhance the adaptive planner, further research into advanced Deep Reinforcement Learning (DRL) techniques could yield a more effective and non-intuitive dosing strategy. Third, the practical MCPS challenges must be addressed. It includes building robust data pipelines, making sure it works with current hospital IT systems and EHRs, and constantly defining the regulatory pathways for approving such a complex AI medical product. Ultimately, this work raises important questions about the ethics of medicine and artificial intelligence. Going forward, it will be important to study the potential biases in algorithms as well as ensuring transparency and interpretability of models. Moreover, the precise and collaborative definition of the role of the clinician in a workflow that is assisted by a fully automated decision support system must be deliberate, and not simply an afterthought, and not must not displace the human element from the ultimate decision-maker in patient care.

This study is limited by its reliance on synthetic patient data and simplified tumor and toxicity models. Although sufficient for proof-of-concept evaluation, real-world clinical heterogeneity may introduce additional challenges. Future work will focus on validation using retrospective

clinical datasets, incorporation of richer mechanistic tumor models, and prospective evaluation within controlled clinical decision-support settings.

References

[1] Abdalla, M., Mohamed, O. Radiomics-enhanced multi-modal deep learning pipelines for stratified cancer prognosis using integrated imaging genomics and clinical data *Int. J. Eng. AI Manag. Decis. Support Policy,* Vol.2, No.2, 2025, pp.45-56.

[2] Arora, S. Transforming AI Decision Support System with Knowledge Graphs & CAG. *IInt. J. Eng. AI Manag. Decis. Support Policy,* Vol.2, No.2, 2025, pp.15-23.

[3] Braman, N., Adoui, M.E., Vulchi, M., Turk, P. Deep learning-based prediction of response to HER2-targeted neoadjuvant chemotherapy from pre-treatment dynamic breast MRI, *arXiv preprint,* arXiv:2001.08570, 2020.

[4] Bouriga, R., Bailleux C., Gal, J., Chamorey, E., et al. Advances and critical aspects in cancer treatment development using digital twins, *Brief. Bioinform.,* Vol.26, No.3, 2025, bbaf237.

[5] Duanmu, H., Huang, P.B.G., Brahmavar, S., Lin, S. Prediction of pathological complete response to neoadjuvant chemotherapy in breast cancer using deep learning with integrative imaging, molecular and demographic data, In: Martel, A.L., et al. *Medical Image Computing and Computer Assisted Intervention* MICCAI, Lecture Notes in Computer Science, Springer, Cham., Vol.12262, 2020, pp.242-252.

[6] Garner, L. How digital twins are shaping the future of cancer care, *MD Anderson Cancer Center,* 2023.

[7] Hormuth, D.A., Jarrett, A.M., Lima, E.A.B.F., McKenna, M.T. Mechanism-based modeling of tumor growth and treatment response constrained by multiparametric imaging data, *JCO Clin. Cancer Inform.,* Vol.3, 2019, pp.1-10.

[8] Komen Foundation. Neoadjuvant Therapies for Breast Cancer, *Susan G. Komen Foundation,* 2023.

[9] Maruf, N.A., Basuhail, A. Breast cancer diagnosis using radiomics-guided DL/ML model systematic review and meta-analysis, *Front. Comput. Sci.,* Vol.7, 2025, Article No.1446270.

[10] Mina, M., Raynaud, F., Tavernari, D., Battistello, E. et al. Conditional selection of genomic alterations dictates cancer evolution and oncogenic dependencies, *Cancer Cell,* Vol.32, No.2, 2017, pp.155-168.

[11] Mohamed, N., Al-Jaroodi, J. A middleware framework to address security issues in integrated multisystem applications, 2019 IEEE International Systems Conference (SysCon), Orlando, FL, USA, 2019, pp.1-6.

[12] Prinzi, F., Orlando, A., Gaglio, S., Vitabile, S. Breast cancer classification through multivariate radiomic time series analysis in DCE-MRI sequences, *Expert Syst. Appl.,* Vol.249, 2024, Article No.123557.

[13] Rawat, P., Saini, R., Kumar, A. Comparative analysis of machine learning techniques for early detection of breast cancer, *Int. J. Comput. Model. Appl.,* Vol.2, No.2, 2025, pp.45-62.

[14] Skarping, I., Larsson, M., Förnvik, D. Analysis of mammograms using artificial intelligence to predict response to neoadjuvant chemotherapy in breast cancer patients: Proof of concept, *Eur. Radiol.,* Vol.32, No.5, 2021, pp.31313141.

[15] Vyas, A., Kumar, K., Sharma, A., Verma, D. et al. Advancing the frontier of artificial intelligence on emerging technologies to redefine cancer diagnosis and care, *Comput. Biol. Med.,* Vol.191, 2025, Article No.110178.

[16] Yang, M., Han, J., Won, H., Chae, S.W. et al. Automated prediction of pathological complete response to neoadjuvant chemotherapy in breast carcinoma using deep learning on pretreatment core needle biopsy samples, *Anticancer Res.* Vol.45, No.11, 2025, pp.5167-5174.

[17] Zeng, H.C., Qiu, S., Zhuang, S.X., Wei, X.L. Deep learning-based predictive model for pathological complete response to neoadjuvant chemotherapy in breast cancer from biopsy pathological images: A multicenter study, *Front. Physiol.,* Vol.15, 2024, Article No.1279982.

[18] Zhou, Y., Shu, X., Wang, F., Xu, H. Prediction of neoadjuvant therapy response in breast cancer based on interpretable artificial intelligence, *Int. J. Surg.*, Vol.112, No.1, 2026, pp.1066-1080.



Xiao Zhou - is an Associate Chief Physician in the Department of Breast Surgery at the Affiliated Cancer Hospital of Guangxi Medical University. She holds a M.Sc. degree in Oncology and a Doctors degree in Surgery from Guangxi Medical University. Her research interests include clinical and basic scientific research related to breast cancer, with expertise in minimally invasive surgery, ductoscopy, sentinel lymph node biopsy, modified radical mastectomy, plastic-conserving mastectomy, breast reconstruction, and systemic therapies.



Longgui Xie - is an Attending Physician in the Department of Breast Surgery at the Affiliated Cancer Hospital of Guangxi Medical University. He holds a Master of Clinical Medicine degree from Guangxi Medical University. His research interests include the metastatic mechanism and drug resistance of breast cancer, as well as breast-conserving surgery, sentinel lymph node biopsy, breast reconstruction, adjuvant/neoadjuvant chemotherapy, endocrine therapy, and targeted therapy.



Pingchuan Li - is a postdoctoral fellow at Guangxi Medical University and a physician in the Department of Breast Surgery at the Affiliated Cancer Hospital of Guangxi Medical University. He holds a Doctor of Clinical Medicine degree from Guangxi Medical University. His research interests include clinical and basic scientific research on breast cancer, and he is proficient in molecular biology techniques.



Jianhui Liu - is an Attending Physician in the Department of Breast Surgery at Guangxi Medical University Cancer Hospital. He is currently a Ph.D. candidate in Oncology at Guangxi Medical University and holds a Master of Medicine in Oncological Surgery from the same university. His research interests include molecular mechanisms of breast cancer progression, drug resistance, targeted therapy, and tumor microenvironment regulation.



Liao Ge-Yi - is a physician in the Department of Breast Surgery at Guangxi Medical University Affiliated Cancer Hospital and a current doctoral student at Guangxi Medical University. He holds a Master of Surgery degree from Guangxi Medical University. His research interests include clinical research related to breast cancer and surgical treatment of breast cancer.



Huawei Yang - is a Professor, Director of the Department of Breast Surgery at the Affiliated Cancer Hospital of Guangxi Medical University, Doctoral Supervisor, and Academic Leader of Breast Surgery. He holds an M.D. degree from Guangxi Medical University. His research interests include clinical and basic scientific research related to breast cancer, with expertise in precise diagnosis, personalized treatment, minimally invasive surgery, ductoscopy, sentinel lymph node biopsy, modified radical mastectomy, plastic-conserving mastectomy, breast reconstruction, and multimodal systemic therapies.