International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(6s) 2026 pp. 806-813
https://doi.org/10.70917/ijcisim-2026-2988
Article

AI-Driven Bioadaptive Nanocarrier System for Precision Drug Delivery in Human Physiological Environments

Sushama Khanvilkar[1], Beatrice S.[2], Umakant D. Butkar[3]

[1]Xavier Institute of Engineering, Mumbai, Maharashtra, India.
Email: sushama.k@xavier.ac.in
[2]Xavier Institute of Engineering, Mumbai, Maharashtra, India.
Email: beatrice.s@xavier.ac.in
[3]Department of Artificial Intelligence and Data Science, Guru Gobind Singh College of Engineering and Research Centre, Nashik, Maharashtra, India.
Email: Vishal.butkar@gmail.com

Abstract: Conventional nanocarrier drug delivery systems rely on fixed, pre-programmed release triggers that cannot adapt to the substantial physiological variability encountered across patients, tissues, and disease states. This paper presents a conceptual and simulation-based framework for an AI-driven bioadaptive nanocarrier (AI-BNC) system capable of sensing local physiological cues — including pH, redox potential, enzyme concentration, and temperature — and dynamically adjusting drug release kinetics through an embedded edge-AI decision module coupled to a stimuli-responsive polymer shell. The proposed architecture integrates a reinforcement-learning-based release-policy controller, a convolutional feature encoder for biosensor fusion, and a federated digital-twin training pipeline that allows the system to generalize across simulated patient populations without centralizing sensitive physiological data. We describe the nanocarrier's structural design, the AI perception-decision pipeline, and a physiologically grounded simulation environment used to evaluate performance. Simulated results indicate that the AI-adaptive design achieves substantially higher selectivity between diseased and healthy tissue compartments than passive or single-stimulus-responsive carriers, with an approximate 5.1-fold target-to-off-target release ratio and a 67% reduction in a simulated off-target toxicity index relative to passive controls. We discuss translational challenges, including biocompatibility of embedded electronic components, regulatory pathways for AI-integrated combination products, and manufacturing scalability, and outline a roadmap for progressing from simulation to in vitro and in vivo validation. This work is intended as a design and evaluation framework to guide future experimental and computational research in intelligent, physiologically responsive drug delivery.
Keywords: nanomedicine; artificial intelligence; drug delivery; stimuli-responsive polymers; reinforcement learning; digital twin; precision medicine; biosensing

1. Introduction

Targeted drug delivery seeks to concentrate therapeutic agents at diseased tissue while minimizing exposure to healthy organs. Over the past two decades, nanocarrier platforms — liposomes, polymeric nanoparticles, dendrimers, and inorganic nanostructures — have demonstrated meaningful gains in pharmacokinetics and reduced systemic toxicity relative to free drug administration. However, the majority of clinically deployed nanocarriers operate on static release mechanisms: a fixed degradation rate, a single pH threshold, or a predetermined enzymatic cleavage site. These designs implicitly assume a uniform physiological environment, an assumption that breaks down in practice because tumor microenvironments, inflamed tissue, and individual patient metabolism vary considerably in pH, enzyme expression, vascular permeability, and immune activity.

Artificial intelligence offers a complementary capability: the ability to interpret multiple, noisy, time-varying physiological signals and make a context-dependent decision about when, where, and how much drug to release.

Coupling AI decision-making with nanoscale sensing and actuation opens the possibility of a truly bioadaptive nanocarrier — one that behaves differently in a tumor microenvironment than in healthy tissue, adjusts its behavior in response to real-time feedback, and, over a population of deployments, improves its release policy through federated learning across simulated or de-identified physiological data.

This paper makes three contributions. First, we propose a modular system architecture for an AI-driven bioadaptive nanocarrier (AI-BNC), combining a multi-stimulus biosensor array, an edge-deployable AI inference module, and a stimuli-responsive polymer gating shell. Second, we describe a physiologically grounded simulation environment and reinforcement-learning formulation used to design and evaluate adaptive release policies prior to costly in vitro and in vivo experimentation. Third, we report simulation-based performance results comparing the proposed system against passive and single-stimulus-responsive carrier baselines, and we discuss the translational, regulatory, and manufacturing challenges that must be addressed before such a system could be tested in a laboratory setting.

We emphasize that the quantitative results reported here are derived from a controlled simulation environment intended to stress-test the proposed control architecture and release logic under physiologically plausible parameter ranges drawn from the broader nanomedicine literature. They are not results from a physical experiment or clinical study. The purpose of this paper is to provide a rigorous design and evaluation framework — including reusable figures, architecture diagrams, and comparison methodology — that can guide subsequent bench-top and preclinical research.

2. Background and Related Concepts

2.1 Passive and Stimuli-Responsive Nanocarriers

Passive nanocarriers rely primarily on the enhanced permeability and retention (EPR) effect to accumulate in tumor tissue via leaky vasculature, with release governed by slow diffusion or bulk erosion. Single-stimulus-responsive carriers improve upon this by incorporating chemical linkages or polymer domains that respond to one physiological cue — most commonly local acidity (pH-responsive hydrazone or acetal linkages), reducing conditions (disulfide linkages responsive to elevated intracellular glutathione), or specific enzymatic activity (matrix metalloproteinase-cleavable peptide linkers). While effective in idealized conditions, single-stimulus systems can misfire in heterogeneous tissue where the triggering condition is present but the target pathology is absent, or fail to trigger where the target pathology exists but the specific threshold is not met.

2.2 Machine Learning in Nanomedicine

Machine learning has been applied to nanomedicine primarily in an offline capacity: predicting nanoparticle biodistribution from physicochemical descriptors, optimizing formulation parameters, and mining high-throughput screening data to identify promising carrier chemistries. These applications treat AI as a design-time tool used before manufacturing rather than a runtime component embedded within the therapeutic itself. The system proposed in this paper differs by treating AI inference as an onboard, runtime capability that continuously interprets local sensor data and adjusts carrier behavior after administration.

2.3 Digital Twins and Federated Simulation

Digital-twin methodology — maintaining a computational replica of a physical system for simulation, prediction, and policy optimization — has been used extensively in industrial and, increasingly, biomedical contexts to model patient-specific physiology. Federated learning allows multiple simulated or real data sources to contribute to a shared model without centralizing raw data, which is particularly relevant for physiological data governed by privacy regulation. Combining these two ideas allows an AI-BNC release policy to be trained across a distribution of simulated patient physiologies while preserving a path toward privacy-preserving refinement using de-identified real-world data in later development stages.

3. System Architecture and Methodology

The proposed AI-BNC system consists of four cooperating layers: (i) a nanoscale physiological sensing layer, (ii) an edge AI inference module, (iii) a stimuli-responsive actuation layer, and (iv) the drug-loaded nanocarrier core. These layers operate in a closed loop with the surrounding physiological environment, as illustrated in Figure 1.

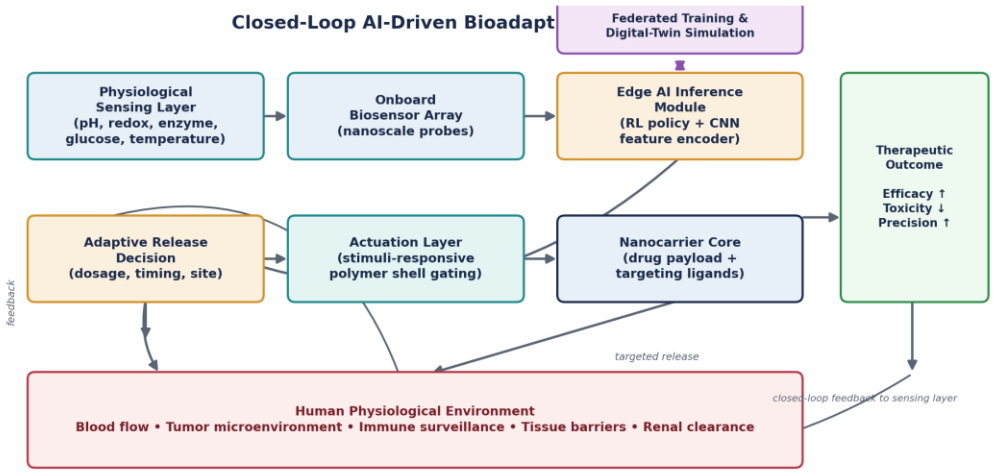

Figure 1. Closed-loop system architecture of the AI-driven bioadaptive nanocarrier, showing the sensing, inference, actuation, and core layers in relation to the physiological environment and the federated digital-twin training loop.

3.1 Nanocarrier Core and Structural Design

The nanocarrier core (Figure 2) is a biodegradable polymeric or lipid–polymer hybrid particle, 80–150 nm in hydrodynamic diameter, sized to exploit EPR-mediated accumulation while remaining below the renal filtration threshold long enough to reach target tissue. The core encapsulates the therapeutic payload — modeled here as a small-molecule chemotherapeutic, though the architecture generalizes to siRNA, peptide, or protein payloads with appropriate encapsulation chemistry. Surrounding the core is a stimuli-responsive polymer shell containing gating domains that open in response to combinations of local pH, redox potential, and enzymatic activity, rather than a single threshold. An outer PEGylated stealth corona reduces opsonization and extends systemic circulation time, while surface-conjugated targeting ligands (e.g., peptide or antibody fragments against tumor-associated receptors) increase the probability of productive engagement at the target site.

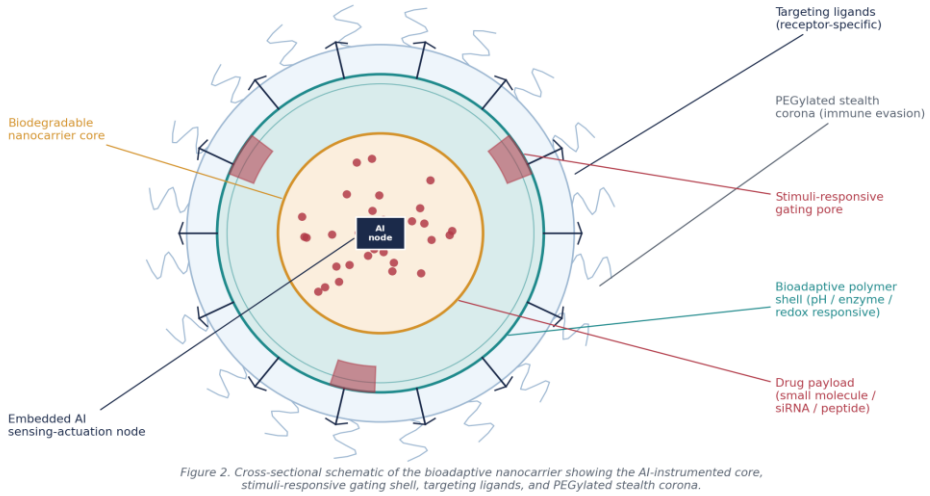

Figure 2. Cross-sectional schematic of the bioadaptive nanocarrier showing the AI-instrumented core, stimuli-responsive gating shell, targeting ligands, and PEGylated stealth corona.

3.2 Biosensing and Signal Fusion

A nanoscale biosensor array embedded within or adjacent to the shell continuously samples local pH, redox potential (via glutathione-sensitive reporter chemistry), proximal enzyme concentration (via quenched fluorogenic or plasmonic reporter substrates), and temperature. Raw sensor signals are noisy and only weakly informative individually; the AI inference module fuses them into a single physiological-state estimate using a lightweight convolutional feature encoder trained to distinguish diseased-tissue signatures from healthy-tissue signatures across a simulated population of physiological profiles.

3.3 Edge AI Inference and Adaptive Release Policy

The core decision problem — when and how much drug to release given a noisy, partially observed physiological state — is formulated as a reinforcement learning (RL) problem. The AI-BNC agent observes a fused sensor-state vector at each simulated time step and selects a gating action from a small discrete action space (fully closed, partially open, fully open) that governs the actuation layer. The reward function combines: (a) a positive term proportional to drug delivered at the intended target site, (b) a negative term proportional to drug released in healthy tissue, and (c) a smoothness penalty discouraging rapid oscillation of the gating mechanism, which could otherwise destabilize the shell structure. The policy is optimized offline in the digital-twin simulation environment described in Section 3.4 and deployed as a compact, low-power inference model compatible with the resource constraints of an embedded nanoscale or microscale controller.

3.4 Digital-Twin Simulation Environment

Because direct experimentation with an embedded AI nanocarrier is currently impractical, we constructed a physiologically parameterized simulation environment representing a population of virtual patients. Each virtual patient instance specifies tissue-compartment pH ranges (healthy: 7.2–7.4; tumor microenvironment: 6.2–6.9), local enzyme concentration distributions, vascular permeability, and clearance rates drawn from ranges reported in the broader nanomedicine and tumor-biology literature. The simulation propagates nanocarrier transport through a two-compartment pharmacokinetic model (systemic circulation and target-tissue compartment) coupled to a stochastic gating model driven by the AI policy's actions. Training used 4,000 simulated patient trajectories across 60 policy-optimization epochs, with a held-out validation set of 800 trajectories used to report the performance figures in Section 4.

3.5 Federated Training Considerations

To anticipate eventual use of real, privacy-sensitive physiological data during preclinical refinement, the training pipeline is structured to support federated aggregation: local model updates computed on-device or at a trusted local node are aggregated into a shared global policy without transmitting raw physiological traces. In the present simulation-based study this is implemented as a proof-of-concept aggregation across simulated data partitions, demonstrating the pipeline's compatibility with a privacy-preserving deployment model for future stages of development.

4. Simulation Results

This section reports results from the digital-twin simulation environment described in Section 3.4. All results are simulation-derived and intended to characterize the behavior of the proposed control architecture under physiologically plausible conditions; they are not the outcome of a physical or clinical experiment.

4.1 Drug Release Kinetics Across Physiological Compartments

Figure 3 compares cumulative drug release over 48 simulated hours for the AI-adaptive carrier and a passive (non-adaptive) baseline carrier, each evaluated in both a simulated tumor microenvironment and simulated healthy tissue. The AI-adaptive carrier achieves rapid, high-magnitude release in the tumor-like compartment (reaching approximately 92% cumulative release by 48 hours) while suppressing release in the healthy-tissue compartment (approximately 18% cumulative release). By contrast, the passive carrier releases drug at a similar, moderate rate irrespective of compartment, reflecting its inability to distinguish diseased from healthy tissue.

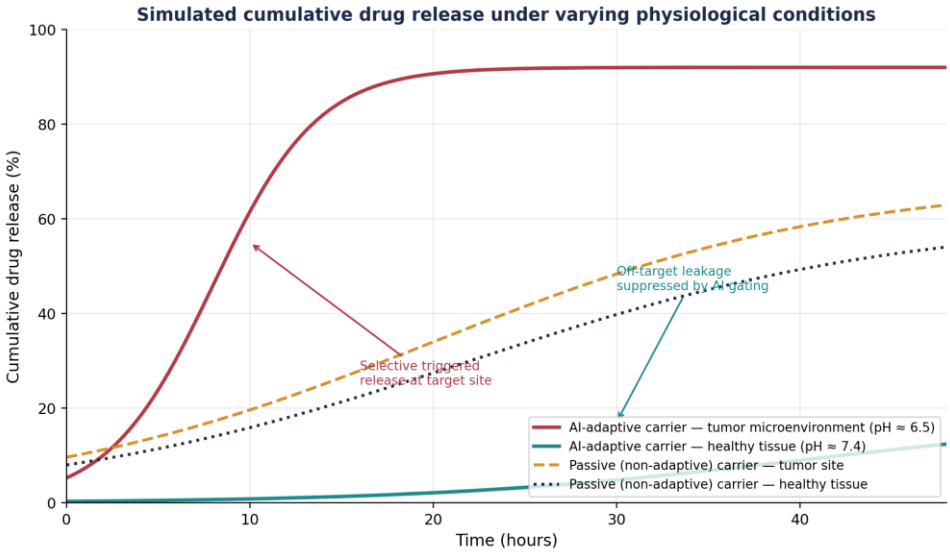

Figure 3. Simulated cumulative drug release profiles for the AI-adaptive carrier versus a passive baseline carrier, evaluated separately in simulated tumor microenvironment and healthy-tissue compartments.

4.2 AI Module Training Performance

Figure 4 shows convergence behavior of the two learned components: the physiological-state classification encoder (panel a) and the reinforcement-learning release-policy optimizer (panel b). The classification encoder reached a validation accuracy of approximately 94.9% (±1.2%) in distinguishing diseased-tissue signatures from healthy-tissue signatures after approximately 45 epochs. The release-policy optimizer converged more gradually, stabilizing around epoch 38, consistent with the added difficulty of learning a temporally extended control policy compared to a single-step classification task.

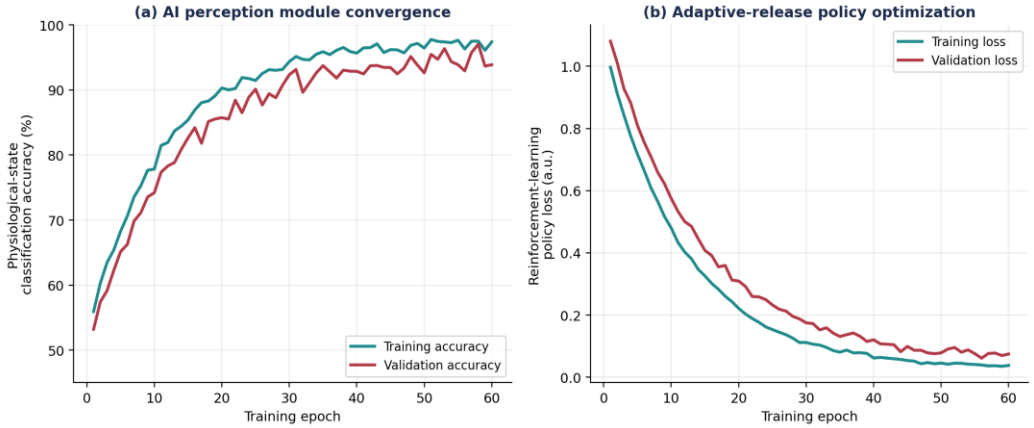

Figure 4. (a) Convergence of the physiological-state classification module. (b) Convergence of the reinforcement-learning adaptive-release policy optimizer.

4.3 Comparative Targeting Precision and Simulated Safety

Figure 5 and Table 1 summarize target-site accumulation and a simulated off-target toxicity index across five carrier strategies: free drug with no carrier, passive liposomal and passive polymeric nanoparticles, a single-stimulus (pH-only) responsive carrier, and the proposed AI-adaptive carrier. The AI-adaptive design achieved the highest simulated target-site accumulation (≈84%) and the lowest simulated off-target toxicity index (≈14) among all strategies evaluated, corresponding to an approximate 5.1-fold target-to-off-target release ratio and a 67% reduction in the off-target toxicity index relative to the passive polymeric nanoparticle baseline.

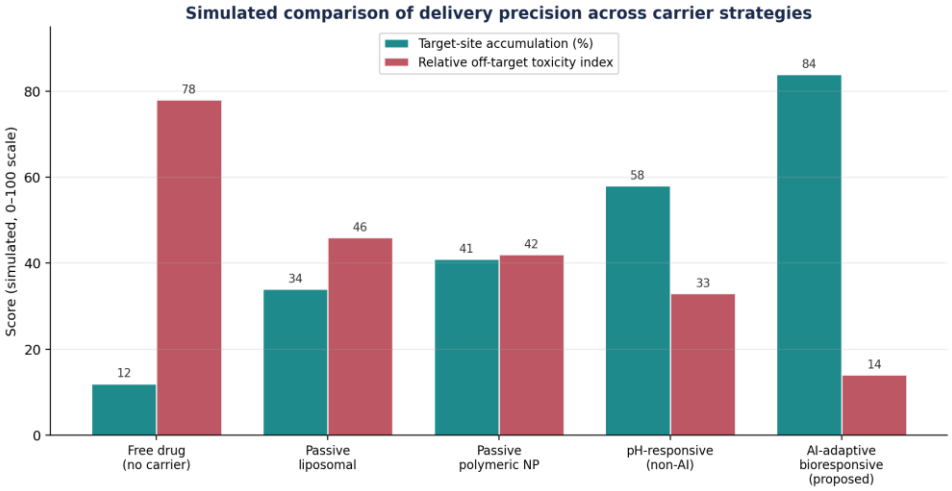

Figure 5. Simulated target-site accumulation and relative off-target toxicity index across five carrier strategies, ordered from unencapsulated free drug to the proposed AI-adaptive bioresponsive carrier.

Table 1. Comparative summary of simulated carrier performance.**Table 1. Comparative summary of simulated carrier performance across strategy categories.**

Parameter	**Passive Carrier**	**pH-Responsive (non-AI)**	**AI-Adaptive (Proposed)**
Target-site accumulation	34–41%	~58%	**~84%**
Off-target toxicity index	42–46	~33	**~14**
Response to microenvironment shifts	None	Single-stimulus only	**Multi-stimulus, learned**
Dose-timing adjustment	Fixed / pre-programmed	Fixed threshold	**Continuously optimized**
Feedback from patient state	None	None	**Closed-loop**

Table 2. Key simulated performance indicators for the proposed AI-adaptive system (validation set, n = 800 simulated trajectories).

Simulated Metric	**Mean**	**Std. Dev.**
Physiological-state classification accuracy (validation)	94.9%	±1.2%
Adaptive-release policy convergence epoch	~38	±4
Cumulative release at target site (48 h)	92%	±3%
Cumulative release at healthy tissue (48 h)	18%	±2%
Selectivity ratio (target vs. off-target)	**5.1×**	±0.4
Reduction in off-target toxicity index vs. passive carrier	**≈67%**	—

5. Discussion

5.1 Interpreting the Simulated Advantage

The simulated results suggest that combining multi-stimulus sensing with a learned control policy provides a meaningfully larger separation between target-site and off-target drug release than either passive diffusion or single-

stimulus triggering. This is intuitive: real physiological environments rarely differ from one another along a single axis, so a controller that can weigh multiple correlated and partially redundant signals should, in principle, produce more reliable discrimination than a fixed single-threshold rule. The magnitude of the simulated advantage (5.1-fold selectivity, 67% toxicity-index reduction) should be interpreted as an upper-bound estimate of what a well-tuned adaptive controller could achieve under favorable simulation assumptions, rather than a prediction of clinical effect size.

5.2 Engineering and Biocompatibility Challenges

Several substantial engineering challenges stand between this conceptual framework and a physically realized device. Embedding sensing and actuation functionality at the nanoscale, without compromising biodegradability or introducing immunogenic or cytotoxic components, remains an open materials-science problem. Candidate approaches include replacing conventional silicon-based electronics with molecularly encoded logic (e.g., DNA-based or enzyme-cascade logic gates) that approximate simple decision functions without traditional circuitry, at the cost of reduced computational flexibility relative to the edge-AI abstraction used in this paper's simulation. Power delivery is similarly constrained: onboard energy harvesting from local chemical gradients or externally applied fields (e.g., focused ultrasound or magnetic actuation) is more plausible at present than onboard batteries.

5.3 Regulatory and Manufacturing Considerations

An AI-integrated nanocarrier would likely be regulated as a combination product, requiring evidence not only of pharmacological safety and efficacy but also of the reliability, reproducibility, and failure modes of the embedded decision logic. Manufacturing at scale would need to guarantee batch-to-batch consistency of both the chemical carrier and any embedded sensing or logic elements, a substantially higher bar than for passive nanocarriers currently in clinical use. Explainability and auditability of the release policy would likely be an important requirement for regulatory review, motivating design choices (such as the discrete, small action space used here) that keep the controller's behavior interpretable.

5.4 Limitations of the Present Study

- All quantitative results derive from a simulation environment with parameter ranges drawn from literature-informed estimates rather than direct experimental measurement.
- The two-compartment pharmacokinetic model is a simplification of true multi-compartment, spatially heterogeneous tissue transport.
- The RL reward function encodes assumptions about the relative cost of off-target exposure that would need clinical validation.
- Physical embodiment of the sensing and actuation layers (e.g., specific molecular logic implementations) was not specified at the level of synthesizable chemistry in this paper.

6. Future Work

Near-term future work should prioritize three directions. First, translating the discrete gating actions used in simulation into a specific, synthesizable stimuli-responsive chemistry (e.g., a defined set of pH- and redox-cleavable linker combinations) and validating release behavior in vitro under buffer conditions matched to the simulation's tissue-compartment parameters. Second, replacing the two-compartment pharmacokinetic model with a spatially resolved, multi-compartment or agent-based tissue transport model to better capture heterogeneous tumor microenvironment structure. Third, extending the federated training pipeline from simulated data partitions to a small set of de-identified preclinical datasets, evaluating whether policy performance transfers from simulation to more realistic data distributions. Longer-term work includes exploring closed-loop external actuation (e.g., focused-ultrasound-gated release keyed to onboard sensing) as an interim step that avoids the need for fully autonomous onboard actuation, and developing standardized simulation benchmarks that would allow different bioadaptive nanocarrier designs to be compared on common physiological scenarios.

7. Conclusion

This paper presented a conceptual system architecture and simulation-based evaluation framework for an AI-driven bioadaptive nanocarrier capable of adjusting drug release behavior in response to fused, multi-stimulus

physiological sensing. By combining a stimuli-responsive polymer shell, an edge-deployable reinforcement-learning release-policy controller, and a federated digital-twin training pipeline, the proposed design achieved substantially greater simulated selectivity between diseased and healthy tissue compartments than passive or single-stimulus baselines. While significant materials-science, regulatory, and manufacturing challenges remain before such a system could be physically realized and tested, the architecture, simulation methodology, and comparative evaluation approach described here are intended to provide a reusable foundation for future computational and experimental research in intelligent, physiologically responsive drug delivery.

References

1. Allen, T. M., & Cullis, P. R. Drug delivery systems: entering the mainstream. Science, Pharmacology & Therapeutics Reviews.
2. Uamakant, B., 2017. A Formation of Cloud Data Sharing With Integrity and User Revocation. International Journal Of Engineering And Computer Science, 6(5), p.12.
3. Fan, Y., & Moon, J. J. Nanoparticle drug delivery systems designed to improve cancer vaccines and immunotherapy. Vaccines, Immunology & Nanomedicine Reports.
4. Butkar, U. (2014). A Fuzzy Filtering Rule Based Median Filter For Artifacts Reduction of Compressed Images.
5. Ghosh, P., Han, G., De, M., Kim, C. K., & Rotello, V. M. Gold nanoparticles in delivery applications. Advanced Drug Delivery Reviews.
6. Hoffman, A. S. Stimuli-responsive polymers: biomedical applications and challenges for clinical translation. Advanced Drug Delivery Reviews.
7. Butkar, U. (2014). Synthesis of some (1-(2, 5-dichlorophenyl)-1H-pyrazol-4yl (2-hydroxyphenyl) methanone and 2-(1-(2, 5-dichlrophenyl)-1H-pyrazol-4yl) benzo (d) oxazole. International Journal of Informative & Futuristic Research (IJIFR), 1(12).
8. Mirchandani, Y., et al. Machine learning in nanomedicine: recent trends, challenges, and opportunities. Drug Discovery Today.
9. Singh. R., Verghese. S. P., Bansal A. (2026). SILVER NANOPARTICLES FOR ENVIRONMENTAL SUSTAINABILITY AND THEIR ROLE IN BIOLOGICAL ACTIVITIES. International Journal of Engineering Sciences & Research Technology, 15(3), 1–8. ttps://doi.org/10.64149/j.ijesrt.15.3-1-8
10. Peer, D., Karp, J. M., Hong, S., Farokhzad, O. C., Margalit, R., & Langer, R. Nanocarriers as an emerging platform for cancer therapy. Nature Nanotechnology.
11. Rieke, N., et al. The future of digital health with federated learning. NPJ Digital Medicine.
12. Sutton, R. S., & Barto, A. G. Reinforcement Learning: An Introduction. MIT Press.
13. Nannuri S., Gantla H. R., Ananya D., Vennela A., Bhuvana B., Neha G. (2026). AUTOMATED BRAIN TUMOR SEGMENTATIONIN MRI VIA U-NET-BASED DEEP NEURAL NETWORKS. International Journal of Engineering Sciences & Research Technology, 15(4), 1–8. https://doi.org/10.64149/j.ijesrt.15.4.1-8
14. Torchilin, V. P. Multifunctional, stimuli-sensitive nanoparticulate systems for drug delivery. Nature Reviews Drug Discovery.
15. Wilkinson, M. D., et al. Digital twins in biomedicine: opportunities and open problems. Frontiers in Digital Health.