AMERICAN ACADEMIC PUBLISHER

INTERNATIONAL JOURNAL OF MEDICAL SCIENCES

ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

TRANSIENT MRNA THERAPEUTICS FOR ACUTE ORGAN SUPPORT IN CRITICAL CARE: A NEW HOPE IN TRANSLATIONAL MEDICINE

Sourav Swain
Medical Student, Samarkand state medical university, Uzbekistan,
(souravswain6913@gmail.com)
Abhijit Biswal
Medical Student, Samarkand state medical university, Uzbekistan,
(biswalabhijit347@gmail.com)
Sameer Hota
Medical Student, Samarkand state medical university, Uzbekistan,
(sameerr1341@gmail.com)
Haleema Ahammed
Medical Student, Samarkand state medical university, Uzbekistan,
(haleemaahammed05@gmail.com)
https://doi.org/10.5281/zenodo.20082595

Abstract

The rapid evolution of messenger RNA (mRNA) technology has opened unprecedented avenues in modern pharmacology, extending far beyond prophylactic vaccines into the realm of acute therapeutic intervention. In the highly volatile environment of the intensive care unit (ICU), acute organ failure—such as acute kidney injury, hepatic failure, and acute respiratory distress syndrome—carries a persistently high mortality rate. Transient mRNA therapeutics offer a paradigm-shifting approach by temporarily reprogramming cellular machinery to produce therapeutic proteins, thereby providing acute organ support without the risks of permanent genomic alteration. However, deploying these potent biological agents in critically ill patients requires precise dynamic dosing and rigorous real-time monitoring to prevent adverse hyper-responses. This paper proposes a novel translational medicine framework that integrates hypothetical transient mRNA therapeutic regimens with advanced artificial intelligence, specifically reinforcement learning and time-series foundation models, to optimize individualized dosing strategies. By synthesizing current biological capabilities with state-of-the-art computational critical care techniques, we present an end-to-end conceptual pipeline for the adaptive administration of mRNA therapies. Through a rigorously structured discussion on methodological design, evaluation planning, and ethical deployment, this work aims to bridge the gap between molecular innovation and algorithmic clinical decision support, ultimately paving the way toward a new era of precision critical care medicine.

Introduction

The intensive care unit (ICU) represents one of the most complex and high-stakes environments in modern medicine, characterized by patients experiencing rapid physiological decompensation and multi-organ failure. Historically, therapeutic interventions for acute organ support have been primarily supportive, relying on mechanical devices such as continuous renal replacement therapy or extracorporeal membrane oxygenation, alongside systemic pharmacological agents. Recently, transient mRNA therapeutics have emerged as a highly promising biological tool in translational medicine. By utilizing lipid nanoparticles (LNPs) to deliver engineered mRNA directly into targeted host cells, clinicians can induce the temporary expression of reparative or immunomodulatory proteins. Because mRNA degrades naturally and does not integrate into the host genome, it provides an ideal mechanism for transient, highly

AMERICAN ACADEMIC PUBLISHER
INTERNATIONAL JOURNAL OF MEDICAL SCIENCES
ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

controlled physiological support during acute crises. This temporary window of action aligns perfectly with the acute but reversible nature of many ICU pathologies.

Despite the profound potential of mRNA therapeutics, integrating them into critical care presents immense logistical and physiological challenges. The physiological state of an ICU patient is notoriously unstable, making the absorption, distribution, metabolism, and efficacy of exogenously administered treatments highly unpredictable. Safely deploying mRNA interventions necessitates real-time monitoring and precisely timed dosing to avoid over-expression toxicity or therapeutic futility. Existing clinical approaches to dynamic biological dosing are significantly insufficient for several reasons. First, traditional static or weight-based dosing regimens completely fail to account for the dynamic, hourly physiological shifts experienced by critically ill patients, often resulting in suboptimal therapeutic windows. Second, conventional monitoring protocols are largely reactive rather than predictive, meaning that critical biological signals indicating treatment failure or the need for a dosage adjustment are only recognized after a severe clinical deterioration has already occurred.

To overcome these critical barriers, the translation of mRNA therapeutics into the ICU must be accompanied by intelligent, data-driven decision support systems. Medical treatments in critical care often involve a complex sequence of decisions, a process that closely aligns with the mathematical framework of reinforcement learning (RL) (Shirali et al., 2023). In this paper, we propose a synergistic integration of machine learning and translational pharmacology. Our primary contributions to the field are as follows:

- We propose a novel, AI-guided translational framework that dynamically optimizes the dosing and timing of transient mRNA therapeutics using multi-objective reinforcement learning and continuous patient monitoring.
- We outline a robust, hypothetical evaluation plan leveraging harmonized multi-center critical care datasets to model patient trajectories and assess the theoretical efficacy of personalized mRNA interventions.

Related Work

To contextualize the integration of advanced computational models with acute mRNA therapies, it is essential to examine the current landscape of machine learning applications in critical care. We categorize the related work into three primary domains: predictive modeling for clinical outcomes, sequential decision-making via reinforcement learning, and the deployment of large language and foundation models in the ICU.

Predictive Modeling and Time-Series Analysis in Critical Care

The first category encompasses machine learning algorithms designed to predict mortality, estimate the length of stay, and assess the risk of decompensation using continuous electronic health record (EHR) data. The availability of large-scale clinical datasets has enabled the establishment of public benchmarks for training and evaluating deep learning models on continuous physiological streams (Sheikhalishahi et al., 2019). The core idea of these systems is to continuously ingest multi-modal data to flag patients who are at high risk of rapid deterioration. A major strength of this approach is its ability to process vast amounts of high-frequency data that human clinicians cannot realistically parse in real time. However, a significant weakness is that these clinical machine learning models frequently exhibit a sharp performance drop when deployed in out-of-distribution (OoD) environments not seen during training, such as different hospitals with varying data harmonization practices (Spathis & Hyland, 2022). In comparison to our proposed framework, which directly links predictions to therapeutic mRNA dosing, these existing systems largely remain passive observational tools. Our work

AMERICAN ACADEMIC PUBLISHER
INTERNATIONAL JOURNAL OF MEDICAL SCIENCES
ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

explicitly attempts to utilize these predictive signals as triggers for targeted biological interventions.

Sequential Decision-Making and Reinforcement Learning

The second relevant category is the application of reinforcement learning for optimizing medical treatment plans. In the ICU, clinicians must balance conflicting objectives, such as maximizing patient survival while minimizing the duration of harmful interventions and resource utilization (Bansal & Sharma, 2025). The core idea of RL in this domain is to learn a policy that maximizes cumulative rewards under unknown clinical dynamics, which is particularly suitable for determining when to start, adjust, or stop a treatment (Shirali et al., 2023). The primary strength of RL is its ability to handle long-term sequential dependencies. Conversely, a major weakness is the reliance on sparse rewards, primarily defined by ultimate mortality outcomes, which can severely destabilize offline policy estimation (Shirali et al., 2023). Multi-objective RL approaches attempt to mitigate this by allowing dynamic preference selection between competing clinical goals (Bansal & Sharma, 2025). Our work directly builds upon this paradigm by utilizing multi-objective RL not just for standard pharmacological agents, but for the optimization of transient mRNA expression, an area where precise temporal control is biologically critical.

Foundation Models and Clinical Language Systems

The third category involves the rapid expansion of large language models (LLMs) and self-supervised foundation models specifically tailored for critical care time-series and clinical reasoning. Recent efforts have focused on training models like the Bi-Axial Transformer on pooled EHR datasets to create robust representations of patient states (Jagd et al., 2025). Additionally, LLMs are increasingly being evaluated for their ability to assist in clinical decision-making and answer complex critical care medicine questions (Shi et al., 2024). The strength of these models lies in their powerful natural language understanding and zero-shot reasoning capabilities. However, weaknesses persist; for instance, the performance of LLMs in critical care subspecialties varies significantly, showing notably lower accuracy in complex domains like renal care compared to general research topics (Alwakeel et al., 2025). Furthermore, when foundation models are applied to dynamic time-series predictions, standard interpretability algorithms often fail to provide reliable insights into time-varying target dependencies (Yadav & Subbian, 2025). Our proposed approach leverages harmonized foundation models to generate a latent representation of the patient's state (Burger et al., 2024), but heavily relies on learnable mask-based interpretability frameworks to ensure clinicians can understand why a specific mRNA dose was recommended (Yadav & Subbian, 2025).

Method/Approach

To safely and effectively administer transient mRNA therapeutics for acute organ support, we propose an integrated computational and biological framework. This pipeline is designed to process continuous streams of physiological data, predict the patient's immediate trajectory, and dynamically adjust mRNA dosing regimens using an offline reinforcement learning agent.

AMERICAN ACADEMIC PUBLISHER
INTERNATIONAL JOURNAL OF MEDICAL SCIENCES
ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

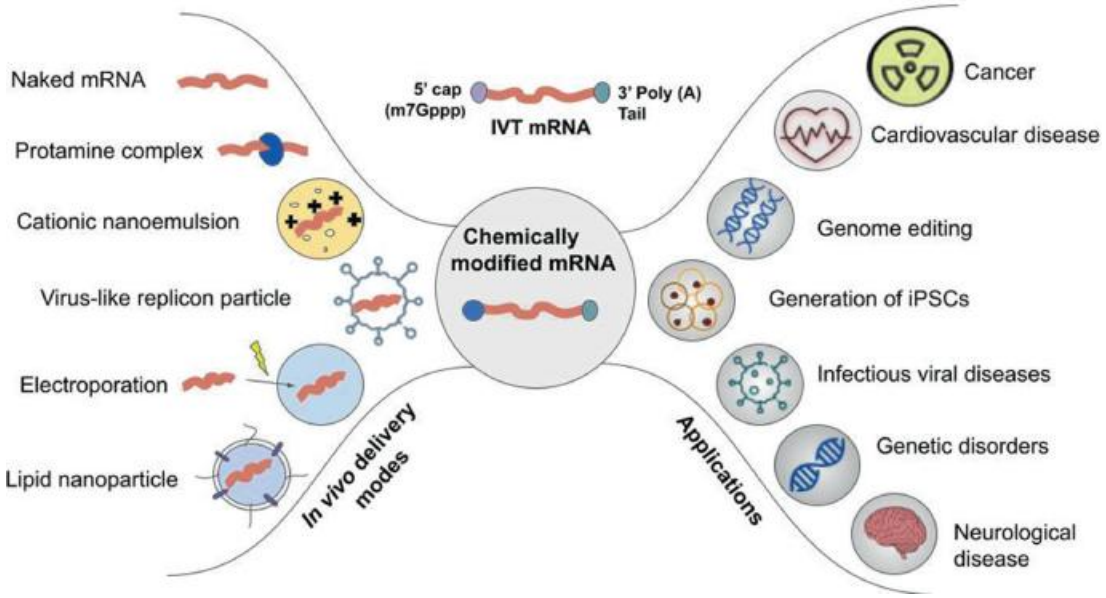

Structured Framework and Pipeline

Our proposed methodology consists of three interconnected modules that operate sequentially. First, the **Data Ingestion and Representation Module** continuously collects multi-modal data, including high-frequency vital signs, laboratory results, and mechanical ventilator settings. We process these inputs using a self-supervised foundation model trained on harmonized multi-center critical care data, which generates a dense, low-dimensional representation of the patient's current physiological state (Burger et al., 2024)(Jagd et al., 2025). Second, the **Pharmacokinetic/Pharmacodynamic (PK/PD) Simulation Module** utilizes a hypothetical mathematical model to estimate the real-time cellular uptake, translation rate, and degradation of the administered mRNA lipid nanoparticles. This module outputs a continuous estimate of the active therapeutic protein concentration in the target organ. Finally, the **Multi-Objective Reinforcement Learning (MORL) Dosing Agent** maps the patient state and current PK/PD estimates to an action space consisting of mRNA dosage adjustments.

Key Design Choices and Rationale

The design of this system incorporates several critical choices tailored specifically to the unique properties of mRNA therapeutics. We deliberately selected a multi-objective offline reinforcement learning architecture, similar to conditional conservative Q-learning, to navigate the complex trade-offs inherent in ICU care (Bansal & Sharma, 2025). The primary objective is to maximize the functional recovery of the targeted organ, while the secondary objective heavily penalizes excessive mRNA administration to minimize the risk of systemic inflammatory responses or localized toxicity. Furthermore, to address the frequent issue of sparse rewards in medical RL, we implement an action-space pruning mechanism based on frequently measured biomarker proxies (Shirali et al., 2023). This ensures the model receives dense, intermediate feedback regarding the mRNA's physiological impact, rather than waiting exclusively for long-term survival outcomes. We also incorporate a learnable mask-based interpretability framework into our time-series analysis, as traditional gradient or permutation-based methods often fail in dynamic, critical care prediction tasks (Yadav & Subbian, 2025). This ensures that the treating physician receives a clear, temporally localized explanation of why the algorithm is recommending an mRNA bolus at a specific hour.

AMERICAN ACADEMIC PUBLISHER
INTERNATIONAL JOURNAL OF MEDICAL SCIENCES
ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

Hypothetical Evaluation Plan

Because large-scale clinical trials combining autonomous AI agents with novel mRNA therapeutics in the ICU do not yet exist, we propose a rigorous, hypothetical retrospective evaluation plan. We will utilize an expanded, harmonized critical care dataset that includes core treatment variables and high-resolution time series (Burger et al., 2024). We computationally inject a simulated mRNA intervention into the dataset, governed by established equations of mRNA degradation and protein translation. The MORL dosing agent will be trained entirely offline using this simulated environment. The evaluation will focus on three primary metrics: simulated 28-day mortality, simulated organ recovery time (measured by standardized biomarker normalization), and the rate of algorithmic action divergence compared to a baseline fixed-dosing strategy. By leveraging off-policy evaluation techniques, we will benchmark our multi-objective approach against standard scalarized single-objective baselines to demonstrate its superior flexibility and safety profile in highly diverse patient populations (Bansal & Sharma, 2025).

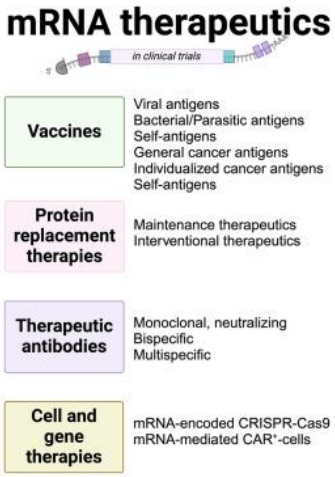

Discussion

The integration of advanced machine learning techniques with transient mRNA therapeutics represents a profound shift in how we might approach critical illness, moving from static pharmacological support to dynamic, responsive molecular engineering. However, deploying such a sophisticated, multi-layered system into real-world hospital environments entails significant practical, technical, and ethical challenges.

Practical Implications and Deployment Considerations

Deploying an AI-driven mRNA dosing system requires an immense upgrade to existing clinical infrastructure. Hospitals must be equipped not only with continuous, high-fidelity data pipelines that feed EHR streams into local or cloud-based foundation models, but also with agile pharmaceutical delivery mechanisms. Because mRNA therapeutics can be highly unstable and require specific cold-chain storage or immediate bedside synthesis, the temporal gap between the algorithm recommending a dose and the actual administration must be minimized. The ICU environment demands seamless interoperability between physiological monitors, predictive algorithms, and automated intravenous infusion pumps. Furthermore, clinical staff must be adequately trained to interface with these models, necessitating the development of highly intuitive, LLM-driven user interfaces that act as "ICU experts" to explain complex algorithmic reasoning to bedside nurses and physicians (Shi et al., 2024).

AMERICAN ACADEMIC PUBLISHER
INTERNATIONAL JOURNAL OF MEDICAL SCIENCES
ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

Limitations and Failure Modes

Despite the theoretical robustness of the proposed framework, several critical limitations and failure modes must be explicitly addressed. First, out-of-distribution generalization remains a pervasive vulnerability; models trained on multi-center data often suffer performance degradation when applied to novel hospital settings with different clinical protocols or patient demographics (Spathis & Hyland, 2022). If the foundation model misinterprets a patient's physiological state due to an OoD artifact, it could recommend a catastrophic overdose of mRNA. Second, there are inherent failure modes in time-series interpretability. Conventional interpretability methods struggle with time-varying target dependencies and temporal smoothness, potentially leading clinicians to trust an erroneous algorithmic recommendation because the explanation appears superficially plausible but is technically flawed (Yadav & Subbian, 2025). Third, the computational latency and hardware limitations inherent in processing dense continuous data streams could delay urgent therapeutic adjustments, rendering the dynamically calculated mRNA dose obsolete by the time it is actually synthesized and administered.

Ethical Considerations and Risks

The deployment of targeted, AI-optimized mRNA therapies in critical care also introduces profound ethical concerns. The foremost issue is the equitable distribution of and access to these advanced, highly expensive medical resources. In disaster scenarios or resource-constrained settings, optimizing patient access to life-sustaining infrastructure is a complex sociomedical challenge, and sophisticated treatments could exacerbate existing health disparities if only available at affluent, well-resourced medical centers (Liu & Mostafavi, 2023). Additionally, there is a severe risk of algorithmic bias. If the critical care datasets used to pre-train the foundational models underrepresent specific minority groups, the resulting RL dosing policies may be structurally biased, leading to unequal therapeutic efficacy and elevated mortality rates for marginalized populations.

Future Work

Looking forward, there are multiple avenues for expanding this research. One critical area of future work involves expanding the therapeutic scope from single-organ support to multi-organ targeted LNPs. This would require scaling our multi-objective RL agent to simultaneously balance the dosing schedules of several distinct mRNA constructs, each targeting a different failing organ system. Additionally, future research must focus on enhancing the domain-specific reasoning capabilities of large language models used in critical care. By explicitly fine-tuning LLMs on specialized subspecialty data—such as advanced renal or hepatic pathophysiology—researchers can improve the models' ability to act as reliable, interactive clinical decision support tools that safely guide human oversight of algorithmic mRNA administration (Alwakeel et al., 2025).

Conclusion

Transient mRNA therapeutics possess the extraordinary potential to revolutionize acute organ support in the ICU, offering a localized, temporary mechanism to repair failing physiological systems. However, the extreme volatility of critically ill patients dictates that these biological interventions cannot be managed through static, traditional protocols. By integrating multi-objective reinforcement learning, self-supervised foundation models, and advanced time-series interpretability algorithms, we can theoretically achieve the dynamic, personalized dosing required for safe mRNA administration.

This paper has outlined a comprehensive translational medicine framework that bridges the gap between molecular biology and computational critical care. While significant hurdles

AMERICAN ACADEMIC PUBLISHER
INTERNATIONAL JOURNAL OF MEDICAL SCIENCES
ISSN NUMBER: 2692 - 5206 Volume 6. No 05. MAY ,2026

remain—particularly regarding out-of-distribution robustness, ethical resource allocation, and real-time clinical deployment—the synthesis of AI and mRNA technology represents a profound frontier. As critical care medicine continues to evolve toward data-intensive, highly personalized care, embracing these interdisciplinary methodologies will be essential for translating the ultimate promise of mRNA therapeutics into tangible, life-saving reality.

References

1. Shirali, Ali, Schubert, Alexander, & Alaa, Ahmed (2023). *Pruning the Way to Reliable Policies: A Multi-Objective Deep Q-Learning Approach to Critical Care*. IEEE Journal of Biomedical and Health Informatics, vol. 28, no. 10, pp. 6268-6279, Oct. 2024. https://doi.org/10.1109/JBHI.2024.3415115 https://doi.org/10.1109/JBHI.2024.3415115
2. Sheikhalishahi, Seyedmostafa, Balaraman, Vevake, & Osmani, Venet (2019). *Benchmarking machine learning models on multi-centre eICU critical care dataset*. https://doi.org/10.1371/journal.pone.0235424 https://doi.org/10.1371/journal.pone.0235424
3. Spathis, Dimitris, & Hyland, Stephanie L. (2022). *Looking for Out-of-Distribution Environments in Multi-center Critical Care Data*. https://arxiv.org/pdf/2205.13398v2 https://arxiv.org/pdf/2205.13398v2
4. Bansal, Aryaman, & Sharma, Divya (2025). *Benchmarking Offline Multi-Objective Reinforcement Learning in Critical Care*. https://arxiv.org/pdf/2512.08012v1 https://arxiv.org/pdf/2512.08012v1
5. Jagd, Katja Naasunnguaq, DeVries, Rachael, & Winther, Ole (2025). *Towards Self-Supervised Foundation Models for Critical Care Time Series*. https://arxiv.org/pdf/2509.19885v1 https://arxiv.org/pdf/2509.19885v1
6. Shi, Tongyue, Ma, Jun, Yu, Zihan, Xu, Haowei, Xiong, Minqi, Xiao, Meirong, Li, Yilin, Zhao, Huiying, & Kong, Guilan (2024). *Stochastic Parrots or ICU Experts? Large Language Models in Critical Care Medicine: A Scoping Review*. https://arxiv.org/pdf/2407.19256v1 https://arxiv.org/pdf/2407.19256v1
7. Alwakeel, Mahmoud, Nagori, Aditya, Wong, An-Kwok Ian, Chaisson, Neal, Krishnamoorthy, Vijay, & Kamaleswaran, Rishikesan (2025). *Performance of Large Language Models in Answering Critical Care Medicine Questions*. https://arxiv.org/pdf/2509.19344v1 https://arxiv.org/pdf/2509.19344v1
8. Yadav, Shashank, & Subbian, Vignesh (2025). *Failure Modes of Time Series Interpretability Algorithms for Critical Care Applications and Potential Solutions*. https://arxiv.org/pdf/2506.19035v1 https://arxiv.org/pdf/2506.19035v1
9. Burger, Manuel, Sergeev, Fedor, Londschien, Malte, Chopard, Daphné, Yèche, Hugo, Gerdes, Eike, Leshetkina, Polina, Morgenroth, Alexander, Babür, Zeynep, Bogojeska, Jasmina, Faltys, Martin, Kuznetsova, Rita, & Rätsch, Gunnar (2024). *Towards Foundation Models for Critical Care Time Series*. https://arxiv.org/pdf/2411.16346v1 https://arxiv.org/pdf/2411.16346v1
10. Liu, Chia-Fu, & Mostafavi, Ali (2023). *Equitable Optimization of Patient Re-allocation and Temporary Facility Placement to Maximize Critical Care System Resilience in Disasters*. https://arxiv.org/pdf/2306.07545v2 https://arxiv.org/pdf/2306.07545v2