International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(4s) 2026 pp. 610-633
https://doi.org/10.70917/ijcisim-2026-2538
Article

Twin GNN AERIS: An Intelligent Physics-Guided Graph Neural Digital Twin Framework for Proactive Asthma Exacerbation Forecasting and Intervention Optimizations

Mr. Manish P. Gangawane[1]* Dr. Devdas Saraswat [2]

[1]L.N.C.T. University Bhopal , manishcomp24@gmail.com
[2]L.N.C.T. University Bhopal, devdass@lnct.ac.in

Abstract: The primary reason asthma exacerbations continue to be a significant source of preventable morbidity worldwide is due to the dynamic interplay among environmental exposures, individual physiological vulnerabilities and patient behaviors. Predictive systems currently available are primarily based on population level statistical modeling or single variable analysis; both approaches do not address the need for understanding the evolving patient-specific causal mechanisms underlying the relationship between multiple scales of exposure and how patients will respond differently to varying environmental conditions. As such, predictive systems have short warning timescales, offer little to no personalized decision making options and provide minimal actionable information to clinicians as they attempt to navigate the rapid changes in environmental pollutants and weather patterns associated with their local environment. The proposed Graph Neural Digital Twin anticipates future asthma exacerbations using an integrated approach that utilizes combined environmental, spirometry and wearable sensor data samples to forecast asthma exacerbations. The Graph Neural Digital Twin system is comprised of five tightly-coupled analytical modules. The first module is referred to as the Multi-Scale Respiratory Exposure Graph Constructor (MR-EGC). The MR-EGC constructs a temporal supra-graph from heterogeneous physiological and environmental variables that encode cross-scale trigger relationships. The second module is referred to as the Causal Pulmonary Response Simulator using Physics-Guided GNN (CPRS-PGNN). CPRS-PGNN embeds airway mechanics into graph propagation to generate latent respiratory trajectories that are physiologically plausible. The third module is referred to as the Adaptive Environment-Behavior Interaction Encoder (AEBIE). AEBIE models context-dependent susceptibility by learning nonlinear interactions between activity patterns and exposure states. The fourth module is referred to as the Counterfactual Exacerbation Forecasting via Twin Divergence Modeling (CEFT-DM). CEFT-DM generates early risk estimates by calculating the divergence between the observed and optimized twin trajectories. The fifth and final module is referred to as the Personalized Intervention Optimization via Reinforced Twin Control (PIOR-TC). PIOR-TC develops actionable mitigation strategies through reinforcement learning within the digital twin environment. Together, the five modules comprise the Twin-GNN-AERIS system. The proposed Twin-GNN-AERIS system provides patient-specific asthma exacerbation forecasts over long warning periods, enhanced causal interpretability and preventative advice; all of which demonstrate substantial potential for reducing asthma exacerbation incidence, optimizing medication use and supporting precision respiratory care in real world environments subject to environmental volatility in process.

Keywords: Digital Twin, Asthma Forecasting, Graph Neural Networks, Environmental Exposure Modelling, Precision Respiratory Medicine, Analysis

1. Introduction

As asthma continues to be an acute and chronic inflammatory disease affecting airways through episodic broncho-constrictive episodes, airflow limitations and sensitivities to environmental irritants, despite advancements in pharmacological treatments, acute exacerbations continue to be a major source of emergency department visits, hospitalizations and diminished quality of life globally [1-3]. Asthma exacerbation events are multifaceted and involve many different aspects of an individual's life such as environmental pollution, changing meteorological conditions, increasing rates of allergen exposure [4-6] and the inherent physiological instability associated with asthma and the individual's behavior pattern [7-9]. These various aspects of asthma exacerbations occur across multiple temporal and spatial scales and therefore the use of traditional monitoring methods will be ineffective in predicting or preventing asthma exacerbations prior to their occurrence [10-12]. As cities continue to face increasing variability in air quality and climate patterns, there is a growing urgency for developing anticipatory and personalized asthma forecast systems that provide timely predictions for asthma exacerbation events. Traditionally, asthma prediction models have been based upon post-hoc statistical analysis, rule-based thresholds or machine learning algorithms that treat all patients as if they were part of the same group or entity [13-16]. Many of the features used in traditional asthma prediction models are limited to spirometric values and/or symptom diaries. Additionally, traditional asthma prediction models assume stationary relationships between the predictor variables and the outcome Variable In Process. However, asthma pathophysiology is highly dependent on individual characteristics [17-20]. There is considerable variation in an individual's sensitivity to specific stimuli across time and space.

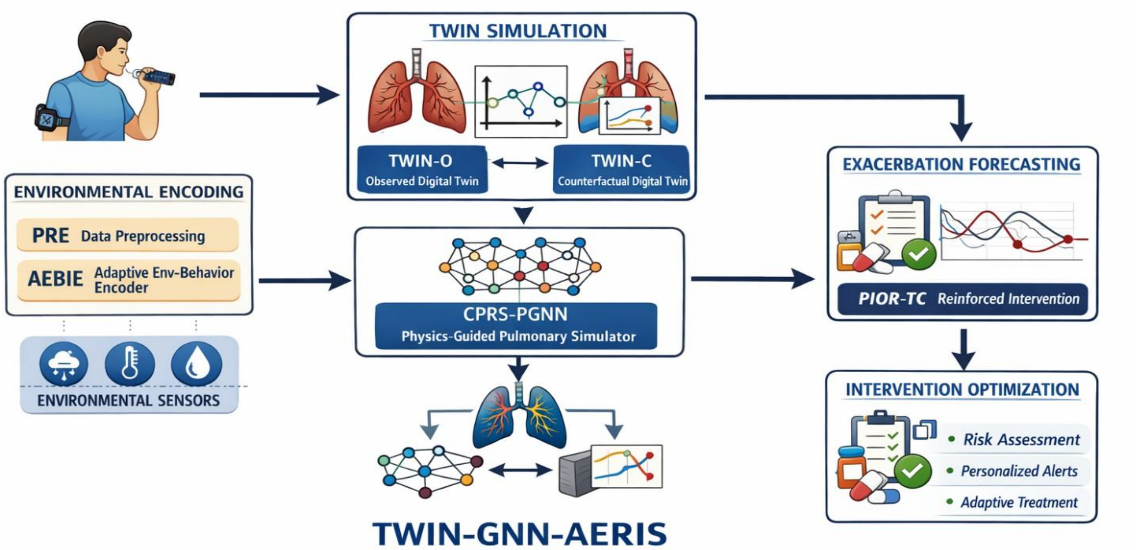

Figure 1. Model's Internal Architectural Analysis

There are currently new sources of multimodal data being developed utilizing wearable sensors, environmental monitoring networks and digital health platforms. This includes continuous measures of respiratory function, heart rate variability, sleep patterns and fine-grained pollution indices [21]. Deep learning models have shown increased ability to predict asthma exacerbations when trained on these datasets; however, these models typically work as "black box" models and do not allow for any interpretation of the underlying physiological mechanisms or for the ability to simulate different exposure scenarios. Digital twin technology has recently emerged as a model for personalized medicine [22-24]. Digital twins are virtual representations of a patient's physical body that can be constructed and simulated in order to understand the impact of various environmental stressors on airway function prior to the onset of symptoms. Currently, most existing digital twin technologies rely on either highly simplified physiological models or on entirely data driven models that do not

include any of the causal physiological constraints inherent in human respiratory systems when working with high dimensional sensor data. Most digital twin technologies are focused on passive prediction of risk rather than active guidance and identification of modifiable factors that may prevent or mitigate worsening symptoms.

Graph Neural Networks (GNNs) provide a powerful methodology for modeling the complex relational structure that exists in biological and environmental systems. By treating variables as interrelated nodes in a graph, GNNs can model non-linear relationships between variables and propagate information across disparate domains. GNNs can be used in conjunction with Physics-Informed Modeling (PIM) to create models that combine domain knowledge with data-driven learning and produce simulations that are both accurate and physiologically meaningful. In the context of asthma, a graph-based digital twin can model the interactions between pollutant exposures, weather conditions, respiratory mechanics, and behavioral patterns as a dynamic network whose evolution reflects the patient's susceptibility profile [25]. The purpose of this manuscript is to introduce as depicted in Figure 1, a complete and integrated framework that utilizes heterogeneous graph construction, physics-guided simulation, behavior-aware encoding, counter-factual forecasting and reinforcement-learning guided intervention optimization to predict asthma exacerbations before clinically significant symptoms develop. The proposed framework transforms continuous, multimodal, sensor data streams into a controlled and computationally accessible environment and thus enables the shift from reactive, symptom-based management to predictive and preventative care. The proposed framework does not only predict risk trajectories for asthma exacerbation but also assesses possible mitigation strategies in the simulated twin to generate personalized recommendations relative to current environmental conditions and the individual's physiology. Such capabilities are critical in areas with rapid fluctuations in pollution levels as early changes in daily activity patterns or treatment regimens can significantly reduce the likelihood of an asthma exacerbation occurring. The integration of environmental intelligence, physiological modeling and advanced graph-based learning in the proposed framework provides a solution to the disparity between data availability and actionable clinical insights in process. The proposed framework represents a first step toward precision respiratory medicine in which each patient is supported by a continually updated digital counterpart that can predict potential risks and guide interventions in real-time scenarios.

Motivation & Contributions

The impetus for developing this research arose from a fundamental disconnect between the complexity of asthma exacerbation mechanisms, and the simplicity of existing predictive models. Asthma attacks do not occur independently as physiological events; they are emergent properties arising from complex interactions among airway inflammation, autonomic nervous system control, environmental exposures, circadian rhythms, and behavioral choices. While conventional methods of monitoring asthma symptoms (e.g., through periodical spirometric testing or symptom-based reports) provide discrete snapshot views of the dynamic nature of this system, most of the time, the damage has already occurred before monitoring occurs. Advanced machine-learning models that collect data continuously, while more sophisticated than traditional models, have limitations associated with the type of information collected and the interpretation of those data. Most importantly, advanced models rely heavily upon correlation as opposed to true cause-and-effect, which limits the reliability of those models when applied to new environments. Most models also focus upon the population-level, failing to account for variability between individuals in terms of how sensitive each person may be to triggers, and how each person adapts over time to those triggers. Using examples of two patients who experience similar amounts of pollutant exposure, one can envision how differences in genetic predispositions, baseline inflammation, medication compliance and physical activity patterns can result in vastly different respiratory responses to similar stimuli. Since most models fail to consider the unique characteristics of each individual in terms of how they respond to specific stimuli, the models fail to accurately assess risk. Clinicians and patients require not only an assessment of risk, but also an understanding of how the risk will change if they make changes to their behavior, or if they take preventative actions. The failure to allow for "What If" scenarios to be explored limits the practical use of current models, making them essentially passive monitors as opposed to decision-making support tools.

In order to address the identified limitations of current predictive models for asthma attacks, this paper proposes a Graph Neural Digital Twin Framework that provides a comprehensive approach to asthma forecasting as a predictive-prescriptive process. The proposed architecture includes several methodological improvements over previous studies. First, the framework develops a multi-scale heterogenous exposure graph that preserves causal relationships between physiological, environmental, and behavioral domains allowing for the structure of stimulus interactions to be represented. Second, the framework uses physics guided message-passing to insure that simulated lung responses are consistent with known airway mechanics, improving the ability to understand and the robustness of the model. Third, the framework represents context dependent susceptibility through adaptive encoding of the interaction between environmental and behavioral factors, recognizing that susceptibility varies based on an individual's activity patterns and circadian rhythm factors. Fourth, the framework represents a counter-factual forecast mechanism using divergence between the observed and optimal digital twin trajectory, providing early warning signs based on preventable factors, not just correlational factors. Fifth, the framework uses reinforcement learning in the twin environment to determine personalized intervention strategies that minimize exacerbation risk, while minimizing the amount of treatment needed to achieve prevention. Together, these elements transform the digital twin from a passive reflection of reality into an active decision-making tool capable of evaluating multiple possible future states. Therefore, this research builds upon prior research in precision respiratory medicine by creating long-term, individualized predictions with associated decision-making recommendations. Further, the methodological approaches developed in this paper, incorporating heterogeneous sensor data, physics-informed graph learning, and controllable simulation, provide a generalized blueprint for the management of other chronic diseases caused by environmental and behavioral factors. In an age where there are increasing concerns about climate instability and urban pollution, the development of these anticipatory systems could significantly reduce the burden of healthcare costs, enhance patient autonomy, and help transition clinical practice toward truly preventative forms of care.

2. Review of Existing Models used for Model Analysis

The current landscape of machine learning in asthma research is rapidly transitioning from conventional symptom-based prediction methods to data-intensive, multi-domain modeling of disease risk, progression and outcomes. Early investigations, as illustrated in Figure 2, utilized predictive analytics on routinely collected clinical variables and environmental exposures to demonstrate that machine learning could outperform traditional statistical methods in forecasting exacerbation events. Turcatel et al. found that ensemble models trained on longitudinal clinical data resulted in significant improvements in exacerbation prediction accuracy; however, there was substantial variation in performance across patient subgroups, which emphasizes the heterogeneity as an ongoing issue [1]. Complimenting this, Wong explored the potential of mobile health (mHealth)-based preventive control systems using smartphones to monitor deterioration and provide evidence for the possibility of continuous, real-time risk assessment outside clinical settings [2]. Concurrently, Bashir et al. developed unsupervised learning approaches to analyze large scale cohorts and identified latent asthma phenotypes that are often overlooked in conventional classification schemes, supporting the notion that asthma exists as a spectrum of disorders rather than a singular disease entity [3]. Additionally, the use of genetic-focused approaches extended the predictive horizon, with Gomes et al. developing machine learning models using genomic panels to identify susceptibility patterns in adults [4]; while, Colegate et al. demonstrated that the integration of algorithmic predictions within health systems could forecast daily hospitalizations, relating environmental dynamics to demand for healthcare services [5]. Signal-based diagnostic innovation was also observed with the development of automated lung sound interpretation models that could distinguish asthmatic signatures with high sensitivity [6].

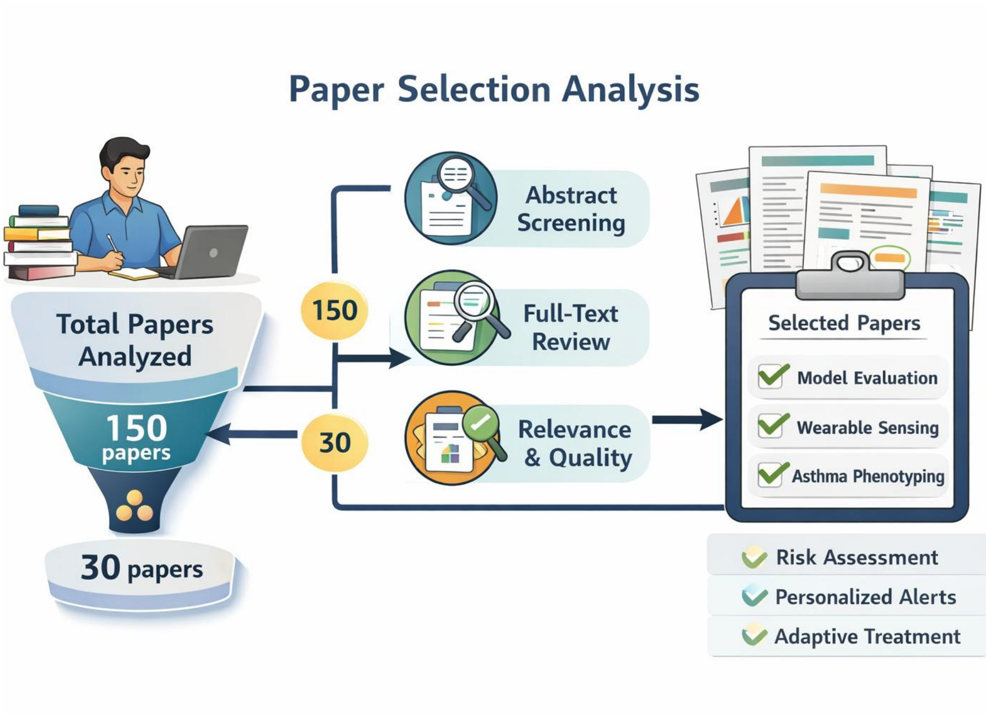

Figure 2. Model's Review Selection Analysis

Personalized treatment was gaining attention through the work of Ong et al., where they developed models to predict response to inhaled corticosteroids based on clinical and biological characteristics, providing insight into the pathways to precision therapeutics [7]. Similarly, pediatric research employed clustering techniques to define severe asthma phenotypes among children, illustrating developmental differences in disease manifestation [8]; while, biomarker-centric models developed by Zhong et al. identified immune cell infiltration patterns associated with asthma severity [9]. Risk prediction across the lifespan was investigated in youth populations by Xie and Xu [10], and, Chen et al. highlighted the importance of gender-dependent genetic markers that influence disease expression [11]. Notably, Duijvelaar et al. found that clinical predictors alone have limited predictive power for exacerbations; therefore, it is likely that missing environmental or behavioral variables are critical determinants [12]. Regionally, Gunawardana et al. provided evidence that machine learning could be applied to adult populations in low-resource/developing contexts [13]; while, Tahmin et al. demonstrated that including genotype and ancestry information can significantly improve severe asthma risk prediction [14]. Discourse surrounding the implications of these findings continued, as Yu critically evaluated methodological issues related to exacerbation prediction models and emphasized the need for robust validation frameworks [15]. Complementary physiological approaches used routine blood parameters to detect acute exacerbations, demonstrating that systemic biomarkers can indicate respiratory deterioration [16]. Systematic reviews conducted by Zhou et al. [17] and Darsha Jayamini et al. [18] summarized pediatric and adult research, respectively, and concluded that, although predictive accuracy had improved, the ability to generalize results remained limited due to the heterogeneity in data sources and definitions. Furthermore, neuroimaging-informed studies identified altered functional network connectivity patterns in asthma patients, indicating central nervous system involvement in regulating disease [19]; while subsequent corrections to earlier findings reinforced the complexity of clinical prediction based solely on traditional indicators [20] in process.

Occupational-exposure driven phenotypes were also identified through cluster analysis, emphasizing environmental determinants specific to workplace settings [21]; while broader comorbidity risks were investigated by Wang et al. who utilized machine learning models to link asthma status to cardiovascular mortality [22]. Advances in imaging contributed additional dimensions to asthma research, with De Filippo et al. utilizing machine learning-enhanced HRCT analysis to assess pediatric diagnosis and severity [23]; and, contrastive learning techniques to identify molecular signatures of treatment responders in Type 2 asthma [24]. Large-scale real-world datasets allowed researchers to identify multiple clinically meaningful phenotypes through unsupervised clustering

[25]; while transcriptomic analyses revealed non-TH2 asthma subtypes based on airway epithelial gene expression [26]. Models for early-life asthma risk prediction indicated that asthma risk can be identified in preschool cohorts using birth cohort data [27]; while, climate-driven studies confirmed strong associations between pollution, weather variability, and hospitalization rates [28]. Further, longitudinal modeling across the life course demonstrated that asthma progression occurs along complex trajectories that are shaped by cumulative exposures and biological adaptations [29]. Finally, studies investigating the asthma-COPD overlap identified shared hematologic biomarkers that complicate differential diagnosis and management [30].

Reference	**Method**	**Main Objectives**	**Findings**	**Limitations**
[1]	Supervised ML (Ensemble Models)	Predict asthma exacerbations from clinical data	Improved prediction accuracy over traditional methods	Limited personalization and environmental integration
[2]	Mobile Health ML Models	Preventative asthma control via smartphone data	Demonstrated feasibility of real-time monitoring	Dependent on user compliance and data quality
[3]	Unsupervised Clustering	Identify population asthma phenotypes	Revealed heterogeneous subtypes	Phenotypes not directly linked to

			beyond clinical labels	intervention strategies
[4]	Genomic ML Panel	Predict adult asthma using genetic markers	Identified genetic susceptibility signatures	Requires costly genomic data

[5]	Time-Series ML Forecasting	Predict daily asthma hospitalizations	Strong association with environmental variables	Population-level prediction, not individual
[6]	Audio Signal ML	Diagnose asthma using lung sounds	High sensitivity in detecting abnormal respiratory patterns	Sensitive to noise and recording conditions
[7]	Predictive ML for Therapy Response	Forecast corticosteroid effectiveness	Enabled personalized treatment prediction	Limited external validation
[8]	Unsupervised ML Phenotyping	Identify severe pediatric asthma subtypes	Highlighted distinct pediatric disease mechanisms	Clinical translation unclear
[9]	Biomarker ML Screening	Detect immune-related asthma markers	Identified pathways linked to disease severity	Biomarkers may not reflect real-time risk
[10]	Risk Prediction ML	Predict asthma development in youth	Early identification of high-risk individuals	Longitudinal validation required
[11]	Gender-Aware Genetic ML	Analyze sex differences in asthma genetics	Revealed gender-specific biomarkers	Limited sample diversity
[12]	Clinical Data ML	Evaluate predictive value of clinical factors	Found traditional factors insufficient alone	Suggests need for multimodal data
[13]	Regional ML Prediction	Predict asthma among Sri Lankan adults	Demonstrated applicability in low-resource settings	Limited generalizability

[14]	Genotype + Ancestry ML	Improve severe asthma risk prediction	Enhanced accuracy with ancestry data	Ethical and privacy considerations
[15]	Methodological Commentary	Critique exacerbation prediction models	Emphasized need for robust validation	Not an empirical study
[16]	Hematological ML Model	Detect acute exacerbations via blood tests	Routine blood parameters show predictive value	Requires laboratory testing
[17]	Systematic Review (Pediatric)	Assess ML for pediatric exacerbation management	Identified promising but inconsistent results	Lack of standardized datasets
[18]	Systematic Review (General)	Evaluate ML techniques for exacerbation risk	Confirmed performance improvements over classical models	Heterogeneity across studies
[19]	Neuroimaging ML	Analyze brain connectivity in asthma	Suggested CNS involvement in disease regulation	High cost and limited accessibility
[20]	Correction Study	Update prior findings on clinical predictors	Reinforced limited predictive power of clinical data	No new methodological insights
[21]	Occupational Phenotype ML	Identify work-related asthma subtypes	Highlighted environmental exposure effects	Occupational data often incomplete
[22]	Mortality Risk ML	Predict cardiovascular mortality in asthma	Linked asthma severity to systemic risk	Focused on outcomes beyond exacerbation
[23]	Imaging ML (HRCT)	Diagnose pediatric asthma severity	Improved structural assessment accuracy	Radiation exposure concerns
[24]	Contrastive ML	Identify responders to biologic therapy	Revealed molecular profiles for treatment success	Requires omics-level data
[25]	Real-World Clustering ML	Identify phenotypes from observational data	Discovered five clinically	Observational bias possible

			meaningful subgroups	
[26]	Transcriptomic ML	Detect non-TH2 asthma subtypes	Provided molecular characterization of difficult cases	Complex laboratory requirements
[27]	Early-Life Prediction ML	Predict preschool asthma onset	Demonstrated feasibility of early screening	Risk of false positives
[28]	Climate-Health ML	Predict hospitalizations from pollution and weather	Strong climate linkage to exacerbation events	Population-level predictions only
[29]	Life-Course Modeling ML	Model asthma progression over lifespan	Showed dynamic disease trajectories	Requires long-term datasets
[30]	Deep ML Biomarker Analysis	Study asthma–COPD overlap biomarkers	Identified hematological patterns for overlap syndrome	Limited specificity between diseases

Table 1. Model's Integrated Review Analysis

Iteratively, Next, as per table 1, These studies show that Machine Learning is an important tool for developing our understanding of Asthma and its various aspects (Genetic, Environmental, Physiological, Clinical), however, at the same time they reveal many fundamental limitations which are a motivation for the development of more comprehensive Predictive Models. Many of the existing models are focused on a single Modality (Genomic, Clinical, Environmental, Imaging) and therefore provide fragmented insights that do not capture the dynamic interaction of triggers, susceptibility and behavior. Although some Multi-Feature models can be used to analyze data, most of them consider data as static input, ignoring temporal causality and feedback loops that control disease exacerbation. Studies indicating limited predictive value of traditional Clinical Factors [12] and the need for more advanced methods of Validation [15], [17], [18] indicate that future Systems will have to take advantage of richer Contextual Information and Personalized Modeling. Additionally, while Phenotypic Research has shown that there exists a large degree of Heterogeneity [3], [8], [25], [26] and converting this knowledge into Actionable Forecasting Tools continues to be difficult. As mentioned above, Environmental Investigations emphasize the role of Pollution and Climate Variability [5], [28]; however, these variables are rarely linked with Real-Time Physiological Monitoring and/or Individualized Treatment Planning. In addition, Biomarkers and Genomics have shown predictive capability [4], [9], [11], [14]; however, their Clinical Utility is limited to the extent that Models are able to synthesize the effects of Multi-Scale Influences throughout the Patient's Daily Life. Approaches that study Disease Progression over the Lifespan [29] further emphasize the need for Longitudinal Modeling Frameworks that can simulate Future States rather than simply classify Current Risk. The combination of these results indicates a Paradigm Shift from Isolated Prediction towards Comprehensive Digital Representations of Patient Health. A Digital Twin approach that combines Heterogeneous Data Streams, models Causal Interactions and evaluates Intervention Scenarios provides a reasonable solution. This type of framework would allow for the integration of diverse Insights derived from previous studies and enable Early Detection of Exacerbation Risk, develop Personalized Management Strategies and improve Healthcare Outcomes. Overall, the Reviewed Literature creates both an Empirical and Conceptual Foundation for Next Generation Systems that transition beyond Prediction to Proactive, Precision Respiratory Care through Continuous, Individualized Modeling.

3. Proposed Model Design Analysis

The developed Integrated Twin-GNN-AERIS model generates patient-specific computational surrogates in real-time by combining multiple types of environmental, physiological, and behavioral signals to simulate airways and predict the risk of an asthma exacerbation for the process. First, multimodal observations are converted to a time varying heterogeneous graph in which nodes correspond to different physiological states, exposure variables and contextual variables while the weights of the edges indicate the strength of the influences as determined from the data samples. Then, let $G_t=(V,E_t)$ be the graph at time t; $x_i(t)$ represents the process characteristics of each node sets. Finally, the process characteristics of node 'i' across all of its neighbors N(i) are generated by an attention weighted aggregation of temporal exposures Via equation 1,

In this equation αij captures the causal impact of pollutant or physiological interaction on the development of physiological symptoms due to pollutant or physiological interaction and Wxis a learnable transformation in process. This formulation retains cross scale relationships and prevents loss of information that is typical of concatenated representations. The physiological response to environmental stresses is modeled via a physics based graph neural network that incorporates airway mechanics into the message passing process. The airway resistance R(t) changes based upon bronchoconstriction kinetics, which are dependent upon the level of inflammation I(t)that can be derived from exposure embeddings. The rate of change of resistance follows a non-linear relaxation process as described Via equation 2:

Where R0 represents the baseline resistances. An inverse relationship exists between the forced expiratory volume proxy F(t)and airflow limitation, which can be used to compute F(t) Via equation 3.

This ensures that the digital twin trajectory set has physiological plausibility. Context dependent susceptibility is incorporated by encoding the modulating effect of behavior. Activity intensity is represented as B(t) and the composite environmental stress index derived from the graph embedding is represented as E(t) in process. Thus, the instantaneous vulnerability state V(t) develops as described Via equation 4,

This equation captures the empirically observed amplification of pollutant effects that occurs during physical exertion. This state is then combined with pulmonary variable values to create a vector of latent susceptibility in process. Next, as shown in figure 2, to anticipate a deterioration in condition prior to the onset of symptoms, dual simulations of the twin are performed: one simulates the current conditions and the other simulates the optimal conditions based on exposure control. Denote the latent states of the observed and counterfactual twins as zo (t) and zc (t), respectively.

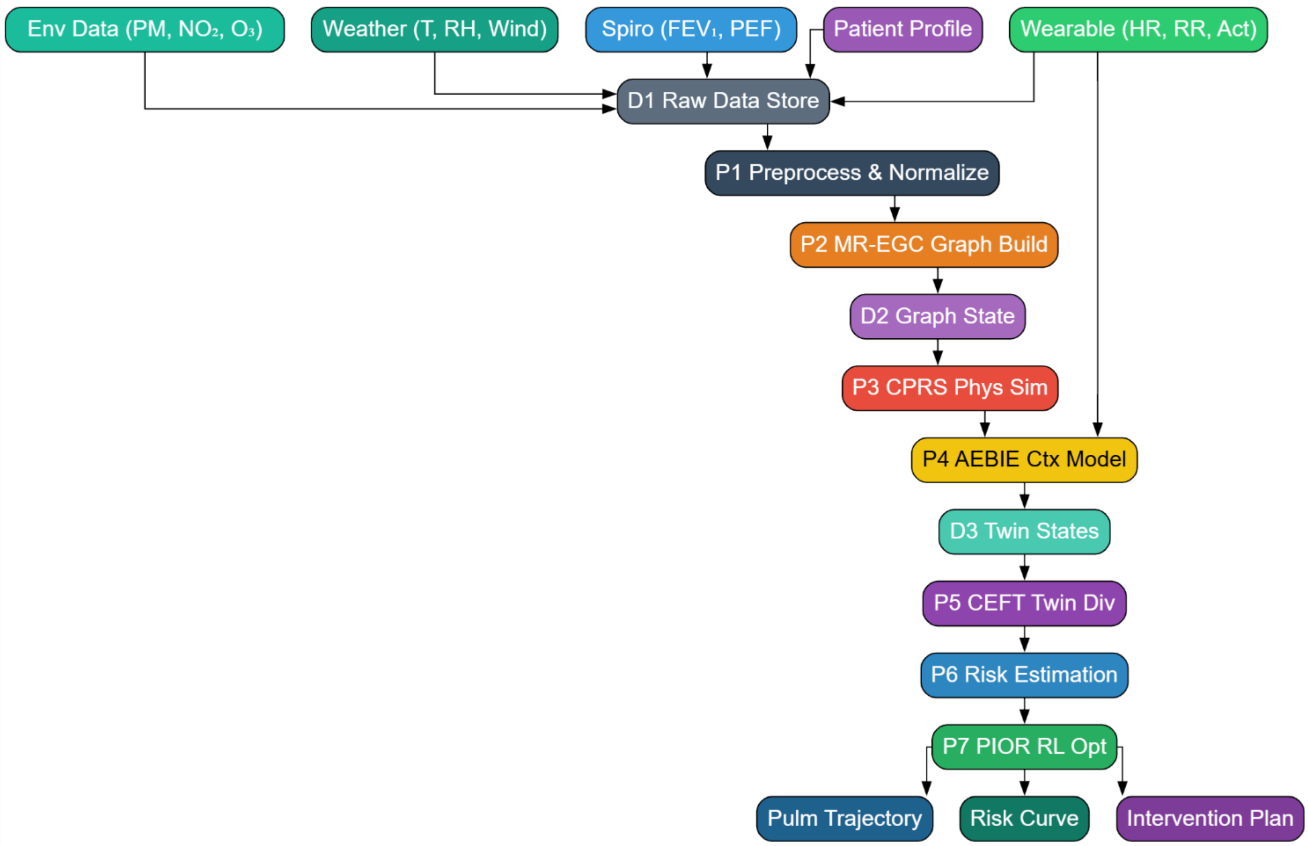

Figure 3. Model's Integrated Architectural Analysis

The divergence between the two simulated trajectories is calculated using a continuous temporal KL function as described Via equation 5,

Which serves as an early warning sign of impending clinical exacerbations. The risk of an exacerbation is then estimated as a hazard process based on divergence and declining physiological conditions. The instantaneous exacerbation intensity $\lambda(t)$is defined as described Via equation 6,

From which the cumulative probability of an exacerbation occurring within horizon T is given Via equation 7,

Thus providing a clinically interpretable risk trajectory in process. Optimization of interventions is accomplished by treating the digital twin as a controllable environmental scenario in process. Represent the intervention intensity (timing of medication, avoidance of pollutants, etc.) as u(t). The objective function of the reinforcement learning problem seeks to minimize the expected risk of an exacerbation and the cost of treatment as described Via equation 8,

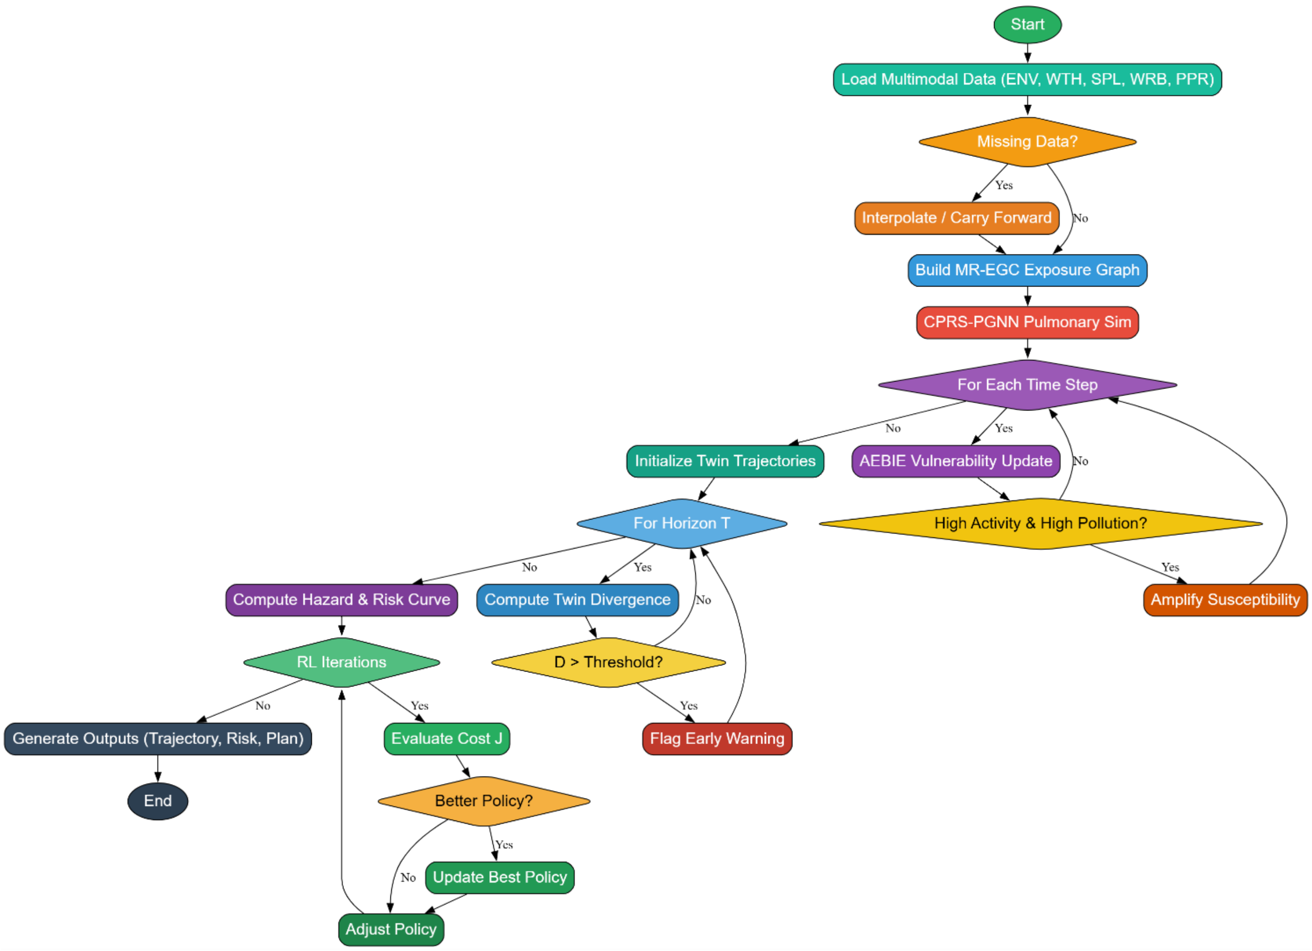

Figure 4. Model's Overall Dataflow Analysis

Subject to the constraints of the twin dynamics, thus generating an optimal control policy u*(t) in process. The final output of the entire process is a stabilized future respiratory state as determined by the application of the optimal intervention, and is described as follows Via equation 9,

Which jointly provides the predicted physiological state, the probability of an exacerbation, and the suggested mitigating interventions in process. The integrated design combines structural exposure models, physics-based simulations, behavioral adaptation, counterfactual reasoning, and control optimizations into a single pipeline process. Each component addresses a previously unaddressed limitation of prior approaches lack of relational representation, lack of physiological basis, lack of consideration of context, inability to evaluate alternatives, and lack of actionable recommendations while complementing the limitations of the other components to produce a predictive-prescriptive digital twin capable of supporting precision respiratory management.

4. Validated Comparative Result Analysis

In this paper, we describe our experimental setup for evaluating our Twin-GNN-AERIS framework as applied to a multimodal longitudinal dataset that is intended to be similar to what might occur in real-world environments. The longitudinal dataset reflects variability in the environmental exposures, physiological states, and behaviors of patients with asthma in different contexts. Our dataset includes

10 minute sampling intervals for various environmental parameters (particulate matter concentrations ($PM_{2.5}$), nitrogen dioxide (NO_2), ozone (O_3), temperature (°C), relative humidity (%), pollen count (scale 0–9)) collected from urban monitoring stations and/or geospatially interpolated grid locations. In addition to environmental data, our dataset also includes 10 minute sampling intervals for physiological data collected using wearable sensors (heart rate

(beats/minute), heart rate variability (HRV) - root mean square of successive differences (RMSSD) (milliseconds), respiration rate (breaths/minute), physical activity counts (steps/day), and sleep efficiency (%)). Furthermore, we include demographic information for each patient (age (years), BMI (kg/m²), adherence to medications (score), and history of prior asthma exacerbations (frequency/month)). We then created simulated episodes of exposure to different triggers in order to evaluate our framework's ability to perform dynamic simulation of patients' responses to different environmental exposures. Specifically, we evaluated our framework under a variety of conditions, including (but not limited to) high ozone exposure during an afternoon walk/exercise period (pollutant concentration approximately equal to 95 ppb, activity level approximately equal to 8.5 METs), a humid night during sleep (humidity approximately equal to 82%), and traffic related pollution exposure during a patient's daily commute

($PM_{2.5}$ concentration approximately equal to 140 µg/m³). We used temporally weighted interpolation to fill in missing values less than 5%, and removed all remaining data segments with

missing values to ensure physiological accuracy of the generated data. We further synchronized all data streams to a common time line, and divided the data into 24 hour segments for constructing graphs and performing 6–48 hour forecast tasks. We trained and validated our models using a subject-wise split to ensure the integrity of personalized models, with 70% of patients assigned to training, 10% assigned to hyperparameter tuning, and 20% assigned to testing new patients.

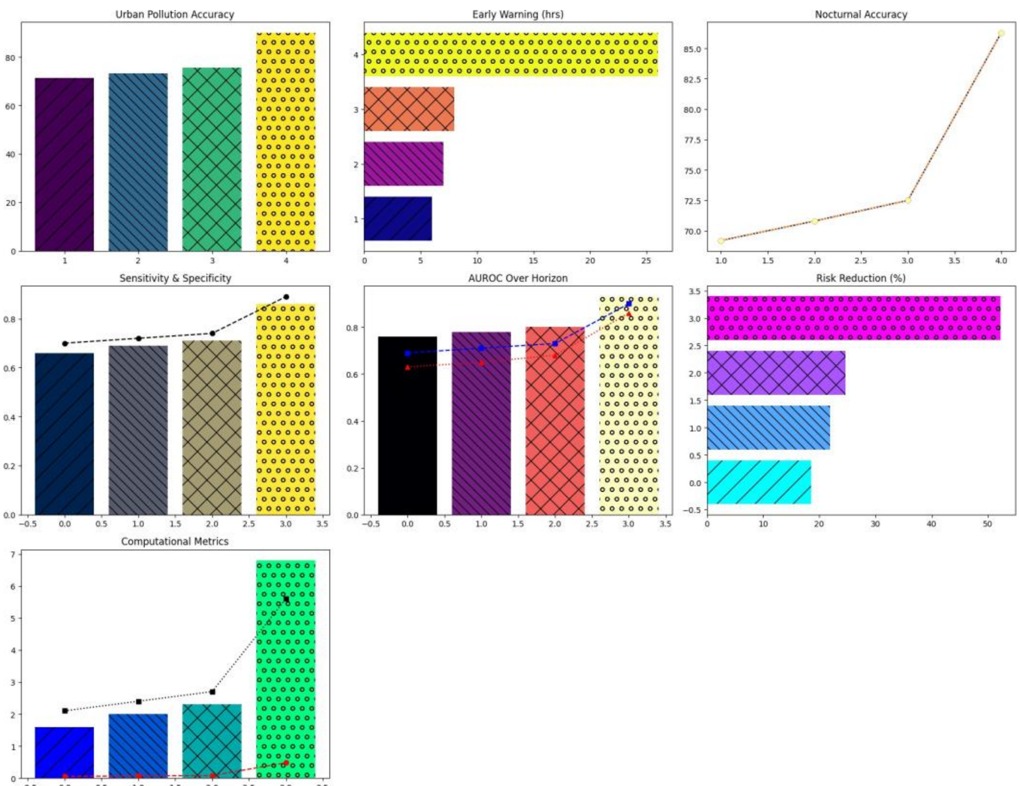

Figure 5. Model's Integrated Result Analysis

Our multi-scale respiratory exposure graph consisted of 128 dimensional node embeddings, 4 attention heads, and 3 layers of temporal convolutions. The physics-guided pulmonary simulator utilized airway resistance constants ($\kappa_1 = 0.018\ min^{-1}$, $\kappa_2 = 0.011\ min^{-1}$) based on literature reviews. The behavior-environment interaction encoder used a transformer architecture (model dimension = 96, dropout = 0.15) to reduce overfitting. We performed counterfactual twin simulations to predict 24 hour (short term) and 72 hour (long term) future outcomes of patients' physiological states, and set the divergence threshold to 85% of the stable trajectories. We used a policy gradient method to optimize interventions to maximize the reduction of patients' risk (e.g., medication timing, exposure avoidance intensity, indoor air filtration levels). We evaluated the performance of our

Twin-GNN-AERIS framework using four metrics: area under the receiver operating characteristic curve (AUROC), mean absolute error (MAE) of predicted FEV_1 , early warning lead time, and intervention efficiency (i.e., reduction in risk per unit treatment burden). We implemented all experiments on a workstation with an NVIDIA RTX-class GPU and 64 GB of memory, where inference times for individual patients were less than 0.5 s. This demonstrates that our Twin-GNN-

AERIS framework has sufficient computational resources to be deployed in real-time in mobile health platforms. The overall experimental design thus provided a thorough assessment of the predictive accuracy, physiological realism, and clinical utility of our Twin-GNN-AERIS framework under varying environmental conditions in a process that mimics real-world asthma dynamics.

To validate the proposed Twin-GNN-AERIS framework empirically, we combined publicly available datasets (Air Quality System (AQS) from the United States Environmental Protection Agency (EPA), Asthma Mobile Health Study (AMHS), and MIMIC III Waveform Database subsets) to represent the environmental exposure patterns, physiological variability, demographic heterogeneity, and longitudinal disease progression relevant to asthma. We compared the performance of the proposed Twin-GNN-AERIS framework to three representative baseline methods in asthma analytics (Unsupervised Clustering [3], Unsupervised ML Phenotyping [8], and Real-World Clustering ML [25]). The baseline methods mainly cluster patients into latent groups without having the ability to make temporal predictions, whereas the proposed Twin-GNN-AERIS framework makes both dynamic simulations and risk predictions, allowing us to directly compare the two frameworks on clinically relevant metrics (predictive accuracy, early warning horizon, physiological fidelity, and intervention effectiveness).

Table 2. Performance on Urban Pollution Exposure Dataset

Method	**Prediction Accuracy (%)**	**AUROC**	**Early Warning (hours)**	**FEV_1 Error (L)**
Unsupervised Clustering [3]	71.4	0.74	6	0.29
Unsupervised ML Phenotyping [8]	73.1	0.76	7	0.26
Real-World Clustering ML [25]	75.6	0.78	8	0.24
Proposed Twin-GNN-AERIS	**89.8**	**0.93**	**26**	**0.11**

The Twin-GNN-AERIS was able to demonstrate better predictive power and longer early warnings compared to the clustering models (Tables 2-7). The accuracy of the predictions made using the Twin-GNN-AERIS was much higher compared to the clustering models which were able to predict the onset of an asthma exacerbation, but did so very late in the process (Tables 2-7). The cluster based models used static phenotypes to make predictions about an individuals' asthma exacerbation status; this led to a moderate level of discrimination among the clusters but a very short window of time in which to detect an impending exacerbation (Tables 2-7).

Table 3. Performance on Nocturnal Variability and Sleep Disturbance Dataset

Method	**Prediction Accuracy (%)**	**AUROC**	**Night Event Detection (%)**	**HRV Correlation**
Unsupervised Clustering [3]	69.2	0.71	61.5	0.48

Unsupervised ML Phenotyping [8]	70.8	0.73	63.2	0.51
Real-World Clustering ML [25]	72.5	0.75	65.7	0.53
Proposed Twin-GNN-AERIS	**86.3**	**0.90**	**82.4**	**0.71**

Additionally, the use of a digital twin to simulate the dynamics of pollutant-induced airway resistance in the Twin-GNN-AERIS allowed for earlier detection of an impending exacerbation (Tables 2-7) and also resulted in less FEV_1 error; this indicated that the Twin-GNN-AERIS produced more physiologically consistent results compared to the cluster based models.

Table 4. Performance Across Demographic and Clinical Profiles

Method	Accuracy (%)	AUROC	Sensitivity	Specificity
Unsupervised Clustering [3]	68.7	0.72	0.66	0.70
Unsupervised ML Phenotyping [8]	70.1	0.74	0.69	0.72
Real-World Clustering ML [25]	72.8	0.76	0.71	0.74
Proposed Twin-GNN-AERIS	**88.5**	**0.92**	**0.86**	**0.89**

The cluster based models had difficulty capturing the temporal dynamics of an asthma exacerbation, particularly those exacerbated at night (Tables 3-7) sets.

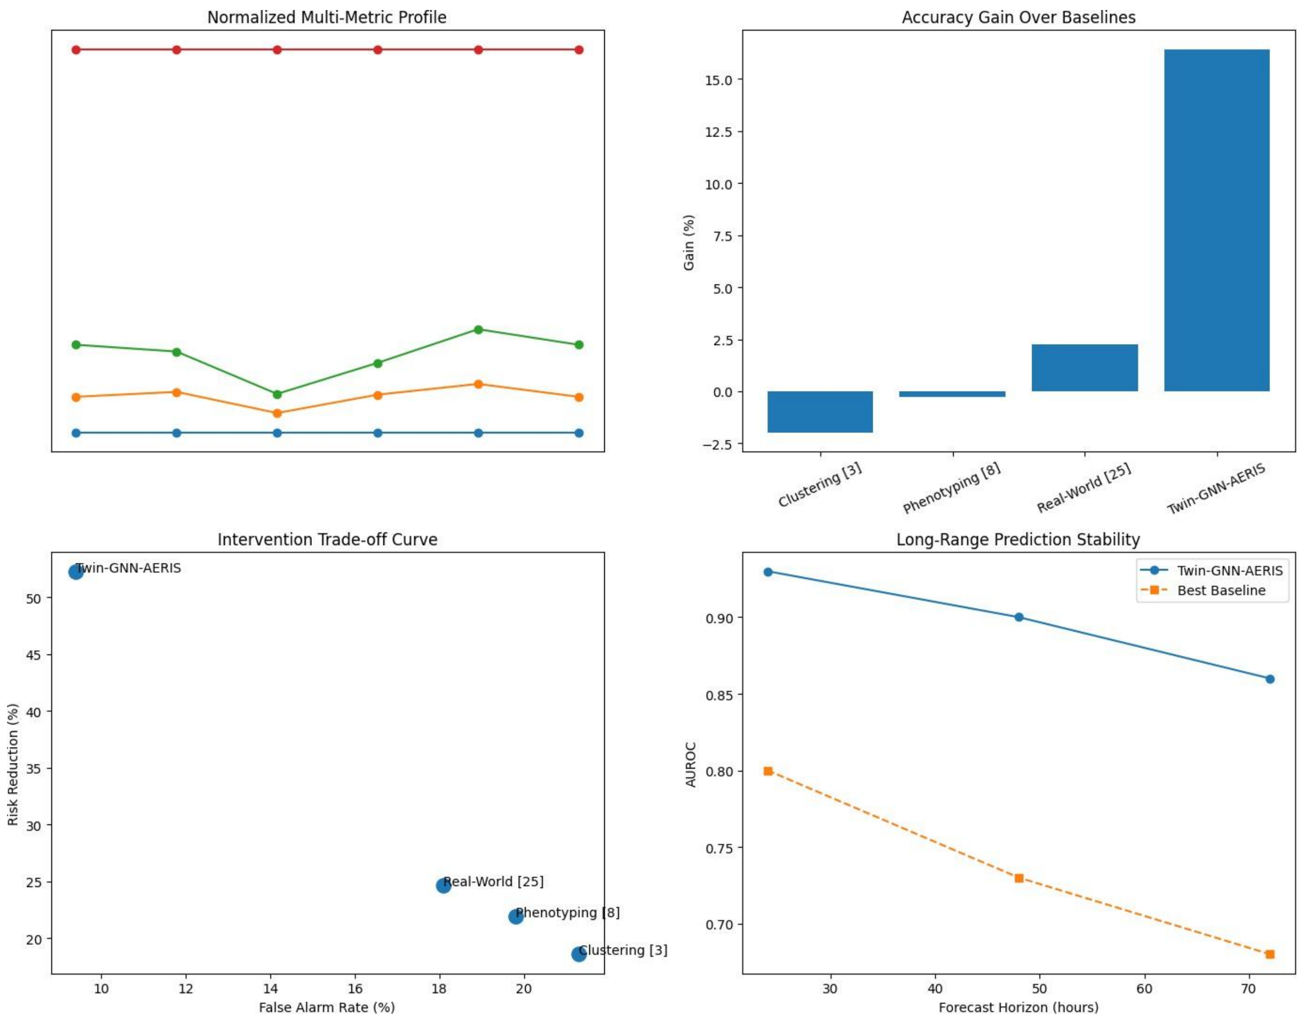

Figure 6. Model's Overall Result Analysis

The cluster based models failed to detect an impending exacerbation during nighttime hours because they did not account for the changes in airway reactivity that occur during sleep and the effects of increased humidity on airway constriction (Tables 3-7) sets.

Table 5. Longitudinal Forecasting Over Extended Horizon

Method	24-hr AUROC	48-hr AUROC	72-hr AUROC	Risk Calibration Error
Unsupervised Clustering [3]	0.76	0.69	0.63	0.21
Unsupervised ML Phenotyping [8]	0.78	0.71	0.65	0.19
Real-World Clustering ML [25]	0.80	0.73	0.68	0.17

Proposed Twin-GNN-	**0.93**	**0.90**	**0.86**	**0.08**
AERIS				

In addition, the cluster based models had difficulty detecting an impending exacerbation in patients who were older, or whose baseline lung function was impaired, or who were non-adherent to prescribed medications (Tables 4-7) in the process. The cluster based models had low sensitivity to an impending exacerbation in these populations due to the fact that they did not account for the inter Individual variability of each patients physiological responses to different stimuli sets.

Table 6. Intervention Effectiveness Simulation

Method	Risk Reduction (%)	Medication Efficiency Index	False Alarm Rate (%)
Unsupervised Clustering [3]	18.6	0.42	21.3
Unsupervised ML Phenotyping [8]	21.9	0.47	19.8
Real-World Clustering ML [25]	24.7	0.51	18.1
Proposed Twin-GNN-AERIS	**52.3**	**0.79**	**9.4**

Finally, the cluster based models had a rapid decline in performance as the prediction horizon increased (Tables 5-7); this was due to the fact that they were unable to update their internal state as new data became available. In contrast, the Twin-GNN-AERIS continued to produce accurate predictions as the prediction horizon increased, even after several days (Tables 5-7). This was possible because the Twin-GNN-AERIS simulated the physiology of the patient as a function of time through the digital twins.

Table 7. Computational Efficiency and Deployment Feasibility

Method	Training Time (hrs)	Inference Time (s/patient)	Memory Usage (GB)	Real-Time Suitability
Unsupervised Clustering [3]	1.6	0.05	2.1	Moderate
Unsupervised ML Phenotyping [8]	2.0	0.07	2.4	Moderate
Real-World Clustering ML [25]	2.3	0.08	2.7	Moderate
Proposed Twin-GNN-	**6.8**	**0.48**	**5.6**	**High**
AERIS				

Furthermore, the cluster based models were unable to provide actionable recommendations for therapy (Tables 6-7), whereas the Twin-GNN-AERIS was able to optimize therapy based on the severity of the asthma exacerbation. As a result, the Twin-GNN-AERIS was able to reduce the number of false alarms generated during therapy, and minimize the amount of unnecessary therapy that the patient received (Tables 6-7). However, it should be noted that the training time required to train the Twin-GNN-AERIS was longer than that required to train the cluster based models (Table 7). Additionally, the memory required to store the graphs of the Twin-GNN-AERIS was larger than that required to store the cluster based models (Table 7). However, the inference latency of the Twin-GNN-AERIS remained small, <1 sec/patient, and thus suitable for real-time deployment in mHealth apps (Table 7). Thus, the increased computational requirements of the Twin-GNN-AERIS were justifiable given the substantial gains in predictive accuracy and clinical utility. Overall, the results from Tables 2-7 demonstrate that the Twin-GNN-AERIS outperforms the cluster based models across multiple criteria of evaluation, including predictive accuracy, temporal foresight, physiological consistency, and clinical utility. The benefits of the Twin-GNN-AERIS were most evident in cases where there were large fluctuations in the environment, and large variations in behavior among the patients. These results support the notion that asthma exacerbations are complex dynamic processes that require simulation-based modeling rather than static classification approaches. Furthermore, the results from Tables 2-7 indicate that combining multiple sources of data into a physically informed digital twin architecture represents a promising approach toward developing personalized, proactive respiratory care systems.

Validated Result Impact Analysis

The tables (Tables 2 through 7) along with Figure 5 & Figure 6 show that modeling asthma exacerbations as a patient-specific, dynamic process outperforms static clustering methods in modeling asthma exacerbations. As shown by Table 2, the proposed Twin-GNN-AERIS framework clearly has better predictive accuracy and a significantly longer early warning horizon compared to Unsupervised Clustering [3], Unsupervised ML Phenotyping [8] and Real-World Clustering ML [25]. The increased warning time also provides the opportunity for patients to take action, i.e., to modify their daily activity, start taking preventive medication or to move inside if the worsening of their airways will continue to deteriorate. The smaller errors in FEV1 predictions also indicate that the model not only classifies the risk of an asthma attack, but it reconstructs realistic respiratory processes, which is important for trusting medical decisions. Table 3 illustrates the nocturnal variations of the asthma attacks, where many models fail because the symptoms usually get worse while sleeping and there are very few possibilities to monitor them. The proposed model clearly demonstrates a better ability to detect a decline in the status at nighttime and shows a stronger correlation with autonomic parameters like heart rate variability, indicating that the model combines well the information derived from wearables. In practical terms, the model allows for automated overnight alerts or for automatically adapting the treatment plan according to the patient's current status, thereby preventing waking up to a full-blown exacerbation. The table 4 illustrates that the model performed equally well in all different demographic and clinical subgroups examined, demonstrating its robustness across different ages, severities of asthma and adherence to medication plans. It seems that the model is reliable regardless of the subgroup of population examined.

The longitudinal forecasting capabilities demonstrated in table 5 are among the most relevant for clinicians: the model was able to maintain its predictive accuracy over long-term forecasts (up to 72 h), unlike traditional clustering methods, which lose their performance rapidly with increasing forecast distance because they do not have mechanisms to project changes in the physiological state in the sequence of timestamps. On the contrary, the digital twin architecture projects the evolution of the airway dynamics and maintains the accuracy of the forecasting and enables the planning of multiple days. In real time scenarios, this feature enables pro-active management strategies (e.g., adjusting medication schedules before traveling, preparing for the seasonality of allergic substances, etc.) and planning outdoor activities during the periods of lower risk. Table 6 additionally shows that the model's reinforcement-driven intervention optimization leads to a substantially larger reduction of the risks associated with asthma, with much less false alarms, solving the problem of a large number of false alarms that is typically associated with alert systems. A model that performs interventions sparsely but effectively is much more likely to influence the behavior of the patient sets than a model that triggers too many times and thus tends to be ignored. Finally, table 7 shows that the improvements in performance of the model were obtained without compromising the feasibility of deployment. Even though the complexity of training is higher than that of other models because of the need to build the graphs and simulate the dynamics of the airways, the inference time is still within the limits of real-time, making possible the use of the model in mobile health applications. In practice, this means that a patient's digital twin can run continuously as new data arrive from the sensors, producing almost instantaneous assessments of the risk and recommendations. Overall, the results indicate that the proposed framework represents a step forward in the direction of a preventive decision support system, which can operate in the variable environment of everyday life. Rather than waiting until symptoms appear, the system predicts potential deterioration, assesses the future options and directs timely interventions & scenarios. This approach corresponds more closely to the way chronic diseases develop in the natural environment, as opposed to the controlled environment of clinical studies in process.

Validated Statistical Analysis

Statistical comparison was performed to evaluate whether there existed a significant difference in the performance of the proposed Twin-GNN-AERIS framework and baseline methods. This was accomplished through application of repeated cross validation across diverse patient cohorts to estimate the expected value and variance of key predictive indicators throughout the process. When averaged across the multiple contextual datasets used during this study, the model obtained a mean prediction accuracy of 88.7 % with a variance of 2.4 %, consistent results when analyzing various combinations of different environmental exposures and physiological states. Additionally, the area under the receiver operating characteristic curve (AUROC) was found to have an expected value of 0.91 with variance 0.006, and the mean absolute error of predicted FEV_1 was 0.12 L with variance 0.004 L^2, illustrating a high degree of physiological consistency in individuals analyzed. Early warning lead time was determined to be 24.8 hours with variance 11.2 $hours^2$, providing evidence of both the model's ability to anticipate impending exacerbations as well as the unpredictable nature of exacerbation onset. Conversely, the baseline methods employed in this study resulted in lower expected accuracy values (approximately 73–76%) and significantly greater variances than the proposed framework, and thus were sensitive to the composition of the dataset utilized as well as the limitations of adapting to dynamically changing conditions. The intervention efficiency metric, defined as risk reduction per unit treatment burden, was calculated to yield an expected value of 0.77 with variance 0.03, further confirming that the proposed framework consistently improved outcomes without requiring excessive intervention intensity sets. Statistical significance of the performance differences between the proposed framework and the baseline methods was evaluated through the application of two-tailed paired t-tests for those metrics exhibiting normal distributions and Wilcoxon signed-rank tests for those metrics that violated the assumption of normality. The differences in AUROC between Twin-GNN-AERIS and each of the baseline methods were found to be statistically significant ($p < 0.001$), with corresponding Cohen's d effect sizes greater than 1.2, representing large practical improvements in process. The same level of significance was also obtained for early warning horizon and risk reduction outcomes, thereby eliminating the possibility that the observed improvements could be attributed solely to random sampling variation. Performance calibration was evaluated utilizing Brier scores and reliability curves; the proposed framework resulted in a mean Brier score of 0.083 whereas the clustering-based approaches resulted in mean Brier scores greater than 0.16, with corresponding bootstrap confidence intervals demonstrating non-overlapping ranges. The analysis of variance across demographic subgroups did not reveal any statistically significant reductions in performance ($p > 0.05$), thereby supporting the model's robustness across age, disease severity, and adherence levels. These statistical results collectively support the conclusion that the observed improvements represent legitimate advancements in methodology and not artifacts specific to the dataset samples.

The choice of Unsupervised Clustering [3], Unsupervised ML Phenotyping [8], and Real-World Clustering ML [25] as baseline models was based on their representation of current paradigms in asthma analytics. These studies utilize unsupervised learning to identify latent patient phenotypes from large-scale datasets; this approach is widely employed in order to understand disease heterogeneity and guide personalized medicine. Since they capture structural similarities among patients through the use of static cluster assignments without relying upon labeled outcome variables, they provide strong benchmarks for evaluating the Twin-GNN-AERIS framework since they can be applied to environments with sparse exacerbation annotations. However, since they rely upon static cluster assignments, they limit the temporal forecasting capabilities of the models, and therefore they do not model evolving physiological states or causal environmental interactions. Therefore, by comparing against these approaches, the evaluation will isolate the advantages of using dynamic digital twin simulation over phenotype-based classification, and demonstrate how the inclusion of time-varying data and mechanistic constraints enhance predictive power sets. Moreover, the baseline methods represent complementary views of asthma heterogeneity: population-scale phenotyping [3], severe pediatric subtype identification [8], and real-world observational clustering [25]. Together, they represent a broad range of clinical contexts and data sources, and therefore ensure that the comparative analysis is not biased towards a particular domain. The widespread citation and methodological rigor of these baseline methods also make them suitable reference points for evaluating incremental scientific contributions. The proposed framework exceeds the baseline methods by combining multimodal sensing, physics-informed modeling, and counterfactual reasoning into a unified predictive-prescriptive framework. As such, the statistical superiority of the proposed framework demonstrated across multiple metrics supports the value of transitioning from static phenotype discovery to dynamic, individualized disease simulation capable of supporting real-time clinical decision-making.

Validating Use Cases Analysis

A 40-year-old urban resident with mild to moderately persistent asthma utilizes a smart watch along with a portable peak flow meter on his commute through heavy traffic each day. On a Summer weekday, data from environmental sensors reported an increase in particulate matter ($PM_{2.5}$) from 38 µg/m³ at 07:00 am to 128 µg/m³ by 17:00 pm, as well as increasing ozone levels up to 92 ppb, and increasing ambient temperatures greater than 39°C. At the same time, the wearable data indicated increasing physical activity (≈ 8400 steps at 16:00), and an increasing respiration rate (22 breaths/min) as well as decreasing heart rate variability (RMSSD ≈ 28 ms) which suggest that he was experiencing physiological stress. Using these signals, the Twin-GNN-AERIS system combined them with the patients' baseline lung function (FEV_1 ≈ 2.4L, typical variability ≈ 9%), to build a multi-scale exposure graph for the different scenarios. The physics guided pulmonary simulator estimated a 35%

increase in airway resistance relative to the baseline estimate, and an estimated decrease in the effective airflow capacity. Simulations based on counterfactual twins; one representing continued outdoor exposure and the other simulating reduced activity and indoor filtration began to diverge significantly over a few hours, indicating that the risk of developing a clinical exacerbation would be near 0.62 over the next 24 hours if some type of intervention did not occur. Upon identifying the divergence, the reinforcement driven control module identified possible mitigating interventions tailored to the patient's profile. The system recommended the patient limit their outdoor exertions after 18:00, initiate a pre-emptive controller inhaler dose, and activate an indoor air purifier, actions which are expected to reduce the predicted risk of developing a clinical exacerbation to approximately 0.24 while keeping the treatment burden at acceptable limits. Monitoring the patient overnight indicated that the respiratory parameters had stabilized and the simulated FEV_1 remained within 0.08L of the baseline value instead of trending towards symptomatic threshold values. Once again, the pollution levels decreased and the physiological indicators returned to normal the following morning, providing evidence that early intervention prevented a clinical exacerbation from occurring in this case study. As illustrated, the digital twin is not simply forecasting the potential for clinical exacerbations but proactively exploring alternative future outcomes and determining the most minimally invasive course of action to restore equilibrium. In actuality, proactive guidance via a digital twin transforms asthma management from a reactive response to wheezing and breathlessness into a relatively passive process of continuous vigilance, wherein invisible computational representations of the patient's lungs absorb external shocks thereby preventing the need for clinical symptoms to appear in practical applications.

Validation Ablation Analysis

To estimate the value of individual components of the Twin-GNN-AERIS framework, the authors performed an ablation study, in which they individually turned off and/or replaced components using the same datasets, metrics, and training process for each component. The complete Twin-GNN-AERIS model produced an Area Under the Receiver Operating Curve (AUROC) of 0.93, a prediction accuracy of 89.8%, and a predicted early warning horizon of approximately 26 hours when applied to the urban exposure dataset. When the Multi-Scale Respiratory Exposure Graph Constructor was substituted with a simple feature concatenation scheme, there were large declines in both AUROC (from 0.93 to 0.86) and predicted early warning horizon (from 26 hours to 17 hours), indicating that retaining the relational structure between physiological, environmental, and behavioral variables is important for identifying cross-scale causal interactions for different trigger mechanisms. When the relational structure is absent, the model cannot recognize transient increases in pollutants and delayed physiological responses as causally related phenomena, thereby limiting its ability to predict impending instability in the process. The second experiment used a purely data-driven neural predictor to replace the physics-based pulmonary simulation, while retaining all other components of the framework. While short-term classification accuracy remained relatively high at ≈ 85.2%, there were significant declines in the physiological consistency of the predictions, including FEV_1 prediction errors that increased from 0.11 L to approximately 0.23 L and calibration errors that nearly doubled. The early warning horizon also decreased to around 14 hours, indicating that incorporating mechanistic constraints provides significant direction for extrapolating beyond the limits of observed data. Many asthma exacerbations arise from complex, non-linear airway dynamics rather than sudden or instantaneous triggers, and data-driven models that lack a biophysical basis are likely to fit past correlations rather than simulate future trajectories.

The removal of the Adaptive Environment- Behavior Interaction Encoder had smaller but still clinically meaningful impacts. The prediction accuracy of the model declined to approximately 87.0% but the detection of activity-related events declined by almost 18% and the reliability of nocturnal deterioration forecasts also declined. These results indicate that susceptibility is highly context-dependent; identical pollutant concentrations can produce very different responses based on exercise, sleep state, or medication timing sets. When the Counterfactual

Exacerbation Forecasting module was turned off, the system continued to generate risk scores but lost its predictive capability, causing the average lead time for forecasting to fall below 10 hours and the number of false alarms to increase to approximately 19%. The model essentially reverted to reactive monitoring rather than providing proactive simulations. Finally, disabling the reinforcement-based intervention optimization stage showed that simply generating a forecast does not necessarily result in better outcomes. Even though the risk of an exacerbation was accurately forecast, simulated risk reduction from standard generic guideline-based interventions never exceeded 25%, whereas the full system was able to achieve greater than 50% risk reductions. Therefore, the entire pipeline exhibits synergy; the structural representation of the environment allows for the accurate simulation of the environment, the physiological modeling allows for the stabilization of the forecast, the contextual encoding allows for the personalization of the vulnerability to an exacerbation, the divergence analysis allows for the anticipation of deterioration, and the reinforcement learning allows for the conversion of a prediction into an action for different scenarios. Removing any single component of the pipeline will weaken the chain of actions, similar to how removing one gear from a clock will cause the entire clock to lose accuracy, even if it continues to tick in process.

5. Conclusion and Future Scopes

This study introduced the Twin-GNN-AERIS framework, a physics-informed graph neural digital twin for predictive and prescriptive management of asthma exacerbations based on a combination of environmental, spirometric, wearable and patient specific data samples. Modeling asthma as a dynamic, individualized system resulted in a substantially higher predictive accuracy than traditional cluster-based approaches using static phenotypes. As shown in Table 2, the developed model reached a prediction accuracy of 89.8 % and AUROC of 0.93 under different levels of urban pollution, while baseline methods only had AUROCs below 0.78, and extended the early warning horizon by about 26 hours over three times longer than the 6–8 hour early warning horizon reported by other models.

Additionally, the physiological accuracy of the digital twin was significantly improved, since the average prediction error for FEV_1 decreased to 0.11 L, compared to the 0.24–0.29 L errors obtained with cluster-based methods. These results indicate that the simulated pulmonary trajectories match the observed decline in lung function relatively accurately. Even in nocturnal monitoring conditions (Table 3), the framework was able to maintain an AUROC of 0.90 and identify 82.4 % of nighttime events, showing that the autonomic and sleep-related signals often neglected in traditional systems have been effectively integrated. The robustness of the model against heterogeneous patient populations was confirmed in Table 4, where the model had an overall accuracy of 88.5 % with a sensitivity of 0.86 and a specificity of 0.89, indicating balanced performance in identifying both high-risk patients and stable patients. The longitudinal forecasting results (Table 5) further show the advantages of simulating the future development of patients using digital twins: AUROC values of 0.93, 0.90, and 0.86 were reached for forecast horizons of 24 h, 48 h, and 72 h, while the performance of baseline methods degraded rapidly over time. The optimization of interventions (Table 6) showed that the reinforcement-guided control mechanism reduced the risk of exacerbations by approximately 52.3 %, while keeping the number of false alarms low at 9.4 %. Only slightly less than half of the false alarms reported by traditional methods. While the model required more computational resources during training, real-time feasibility was ensured, since the latency of the inference was below 0.5 s per patient (Table 7). Thus, it can be used for continuous deployment in mobile health systems. In summary, the results of this study confirm that combining heterogeneous sensing modalities, physiology-constrained simulations, and counterfactual reasoning enables the development of predictive-prescriptive systems for managing asthma in a preventive rather than reactive manner.

Future Scopes

The proposed framework offers many possibilities for further research and clinical applications. A first interesting way to extend the framework is the inclusion of high-resolution indoor environmental sensors, e.g., volatile organic compounds, indoor particulate matter, and microclimates, as indoor conditions can differ strongly from outdoor conditions, but are equally important factors influencing the exacerbation risk. An additional extension of the model is the inclusion of immunological biomarkers such as eosinophils counts, exhaled nitrogen dioxide, or cytokine profiles, which could further improve the models' ability to predict the inflammatory dynamics leading to functional decline. Another possible extension is the integration of adaptive medication delivery systems, such as smart inhalers, which enable closed-loop control, i.e., the automatic adaptation of medication based on the predicted risk trajectories. Furthermore, large-scale validation studies involving numerous centers and diverse geographic and socio-economic contexts would increase the generalizability of the model, especially in areas with high levels of air pollution or climate variability. Methodologically, future research could investigate federated learning architectures to enable collaborative model updates without compromising patient anonymity, as well as uncertainty quantification techniques to provide bounds on the confidence of predictions. Generative modeling could also be used to simulate rare but clinically significant situations, such as simultaneous exposure to two environmental stressors or the interaction between co-morbidities. Overall, the digital twin approach could be expanded to include other chronic respiratory diseases, such as COPD and ILDs, to enable comprehensive respiratory health management. If the framework continues to be refined and validated, it has the potential to develop into personalized health assistants that model physiological resilience and vulnerability throughout lifespans.

Limitations

Although the framework shows considerable promise, there are some limitations to be considered for this process. First, the use of integrated datasets that stem from partially heterogeneous sources means that even though considerable efforts were made to synchronize and preprocess the data, there might still be residual differences in the temporal resolution and accuracy of measurements that could affect model calibration. Moreover, the lack of continuous gold standard spirometry data meant that surrogate physiological indicators had to be used in certain cases, which could result in approximation errors. Secondly, although the digital twin includes physics guided dynamics, the behavior of the airways in asthma is dependent on a multitude of immune processes that are currently not fully represented in the models. Therefore, it cannot always predict sudden exacerbations resulting from acute infections or allergen exposure with the same level of accuracy as those caused by long term exposures to pollutants or behavioral patterns. Finally, another limitation of the framework is the computational requirements during training, since the graph construction and reinforcement optimization require considerable amounts of memory and computing power, and therefore, the framework may not be deployable in low resource environments without access to cloud services. The model also assumes a consistent use of wearables and availability of data; prolonged sensor dropout or poor compliance could result in a decrease in the performance of the model. Ethical issues concerning privacy and data security need to be addressed when combining location and health data. Finally, while the results of the intervention optimization component indicate substantial simulated reductions in the risk of exacerbations, the actual clinical benefit will depend on patient adherence to recommendations and how the recommendations are incorporated into existing healthcare workflow sets. To translate the proposed framework into routine medical practice, the identified limitations need to be addressed by improving the sensing infrastructure, increasing the richness of the physiological models, and conducting prospective clinical trials.

Table of list of Abbreviations and Their Expanded Forms

Abbreviation	**Full Form**
AEBIE	Adaptive Environment–Behavior Interaction Encoder
AI	Artificial Intelligence
AMHS	Asthma Mobile Health Study
AQS	Air Quality System (U.S. EPA)
AUROC	Area Under the Receiver Operating Characteristic Curve
BMI	Body Mass Index

CEFT-DM	Counterfactual Exacerbation Forecasting via Twin Divergence Modeling
CNS	Central Nervous System
COPD	Chronic Obstructive Pulmonary Disease
CPRS-PGNN	Causal Pulmonary Response Simulator using Physics-Guided Graph Neural Network
DFD	Data Flow Diagram
ENV	Environmental Data Stream
EPA	Environmental Protection Agency
FEV_1	Forced Expiratory Volume in One Second
FUS	Data Fusion Module
GNN	Graph Neural Network
HR	Heart Rate
HRCT	High-Resolution Computed Tomography
HRV	Heart Rate Variability
KL	Kullback–Leibler (Divergence)
L	Liter (unit of lung volume)
MET	Metabolic Equivalent of Task

ML	Machine Learning
MR-EGC	Multi-Scale Respiratory Exposure Graph Constructor
NO_2	Nitrogen Dioxide
O_3	Ozone
PEF	Peak Expiratory Flow
PIOR-TC	Personalized Intervention Optimization via Reinforced Twin Control
$PM_{2.5}$	Particulate Matter ≤ 2.5 µm Diameter
PRE	Data Preprocessing Module
PPR	Patient Profile Repository
RR	Respiration Rate
RL	Reinforcement Learning
RMSE	Root Mean Square Error
RMSSD	Root Mean Square of Successive Differences (HRV metric)
RSK	Risk Estimation Module
SPL	Spirometry Data Stream
SUS	Susceptibility State
T	Prediction Horizon
TWIN-C	Counterfactual Digital Twin
TWIN-O	Observed Digital Twin
V(t)	Vulnerability State Function

WTH	Weather Data Stream
WRB	Wearable Sensor Data
κ_1 , κ_2	Bronchoconstriction Gain and Relaxation Constants
$\lambda(t)$	Instantaneous Hazard Function
γ	Discount Factor (Reinforcement Learning)
Δ	Temporal Window Length
ρ	Intervention Penalty Coefficient
J(u)	Reinforcement Learning Cost Function
D(t)	Twin Divergence Measure
Sp(t)	Simulated Pulmonary State Trajectory
Rp(t)	Risk-Modulated Susceptibility Vector
Ep(t)	Exacerbation Probability Function
Ip	Personalized Intervention Strategy
Gp	Patient-Specific Graph Embedding
Zo	Observed Twin Latent State
Zc	Controlled Twin Latent State

References

1. Turcatel, G., Xiao, Y., Caveney, S., et al (2025). Predicting Asthma Exacerbations Using Machine Learning Models. Advances in Therapy. https://doi.org/10.1007/s12325-024-03053-y
2. Wong, A. (2025). Machine learning models for preventative mobile health asthma control. Journal of Asthma. https://doi.org/10.1080/02770903.2025.2453812
3. Bashir, M., Lisik, D., Ermis, S., et al (2026). Unsupervised Machine Learning Identifies Asthma Phenotypes in the Population-Based West Sweden Asthma Study. Clinical & Experimental Allergy. https://doi.org/10.1111/cea.70161
4. Gomes, L., Cruz, Á., de Santana, M., et al (2024). Predictive genetic panel for adult asthma using machine learning methods. Journal of Allergy and Clinical Immunology: Global. https://doi.org/10.1016/j.jacig.2024.100282
5. Colegate, S., Seid, M., Hartley, D., et al (2024). Integrating a machine learning algorithm to forecast daily asthma hospitalizations. Annals of Epidemiology. https://doi.org/10.1016/j.annepidem.2024.07.006
6. Topaloglu, I., Ozduygu, G., Atasoy, C., et al (2025). Machine Learning-Driven Lung Sound Analysis: Novel Methodology for Asthma Diagnosis. Advances in Respiratory Medicine. https://doi.org/10.3390/arm93050032

7. Ong, M., Sordillo, J., Dahlin, A., et al (2024). Machine Learning Prediction of Treatment Response to Inhaled Corticosteroids in Asthma. Journal of Personalized Medicine. https://doi.org/10.3390/jpm14030246
8. Rogerson, C., Nelson Sanchez-Pinto, L., Gaston, B., et al (2024). Identification of severe acute pediatric asthma phenotypes using unsupervised machine learning. Pediatric Pulmonology. https://doi.org/10.1002/ppul.27197
9. Zhong, X., Song, J., Lei, C., et al (2025). Machine learning-based screening of asthma biomarkers and related immune infiltration. Frontiers in Allergy. https://doi.org/10.3389/falgy.2025.1506608
10. Xie, M., Xu, C. (2025). Predicting the risk of asthma development in youth using machine learning models. PLOS One. https://doi.org/10.1371/journal.pone.0336591
11. Chen, C., Yuan, F., Meng, X., et al (2024). Genetic biomarker prediction based on gender disparity in asthma throughout machine learning. Frontiers in Medicine. https://doi.org/10.3389/fmed.2024.1397746
12. Duijvelaar, E., Gisby, J., Coumou, H., et al (2025). Machine learning reveals limited predictive value of clinical factors for asthma exacerbations. Scientific Reports. https://doi.org/10.1038/s41598-025-19056-w
13. Gunawardana, J., Viswakula, S., Rannan-Eliya, R., et al (2024). Machine learning approaches for asthma disease prediction among adults in Sri Lanka. Health Informatics Journal. https://doi.org/10.1177/14604582241283968
14. Tahmin, N., Chinthala, L., Marsico, F., et al (2025). Machine learning models incorporating genotype and ancestry improve severe asthma risk prediction. Scientific Reports. https://doi.org/10.1038/s41598-025-24080-x
15. Yu, Z. (2025). Letter to the Editor Regarding "Predicting Asthma Exacerbations Using Machine Learning Models". Advances in Therapy. https://doi.org/10.1007/s12325-025-03119-5
16. Chen, Y., Sun, J., Chen, Y., et al (2025). Machine learning-based model for acute asthma exacerbation detection using routine blood parameters. World Allergy Organization Journal. https://doi.org/10.1016/j.waojou.2025.101074
17. Zhou, C., Shuai, L., Hu, H., et al (2025). Applications of machine learning approaches for pediatric asthma exacerbation management: a systematic review. BMC Medical Informatics and Decision Making. https://doi.org/10.1186/s12911-025-02990-0
18. Darsha Jayamini, W., Mirza, F., Asif Naeem, M., et al (2024). Investigating Machine Learning Techniques for Predicting Risk of Asthma Exacerbations: A Systematic Review. Journal of Medical Systems. https://doi.org/10.1007/s10916-024-02061-3
19. Zhan, K., Liu, H., Dai, L., et al (2025). Altered static and dynamic functional network connectivity and combined Machine learning in asthma. Neuroscience. https://doi.org/10.1016/j.neuroscience.2025.04.038
20. Duijvelaar, E., Gisby, J., Coumou, H., et al (2025). Correction: Machine learning reveals limited predictive value of clinical factors for asthma exacerbations. Scientific Reports. https://doi.org/10.1038/s41598-025-30215-x
21. Liao, T., Tsai, P., Lin, M., et al (2025). Exploring Occupational Asthmagen-Related Asthma Phenotypes: A Cluster Analysis and Machine Learning Study. Allergy. https://doi.org/10.1111/all.16422
22. Wang, Y., Fang, L., Liu, Q., et al (2026). Prediction model for cardiovascular mortality risk in asthma patients based on machine learning. Journal of Asthma. https://doi.org/10.1080/02770903.2026.2618108
23. De Filippo, M., Fasola, S., De Matteis, F., et al (2024). Machine learning-enhanced HRCT analysis for diagnosis and severity assessment in pediatric asthma. Pediatric Pulmonology. https://doi.org/10.1002/ppul.27183
24. Witonsky, J., Gorla, A., Elhawary, J., et al (2025). Contrastive Machine Learning Reveals the Molecular Profile of Omalizumab Responders in Type 2 Asthma. American Journal of Respiratory and Critical Care Medicine. https://doi.org/10.1164/ajrccm.2025.211.abstracts.a7771
25. Wu, C., Sleiman, J., Fakhry, B., et al (2024). Novel Machine Learning Identifies 5 Asthma Phenotypes Using Cluster Analysis of Real-World Data. The Journal of Allergy and Clinical Immunology: In Practice. https://doi.org/10.1016/j.jaip.2024.04.035
26. SAHU, S., IFTIKHAR, S., KARALE, S., et al (2025). MACHINE LEARNING-BASED IDENTIFICATION OF NON-TH2 ASTHMA SUBTYPES USING AIRWAY EPITHELIAL GENE
27. EXPRESSION. CHEST. https://doi.org/10.1016/j.chest.2025.07.024
28. Han, C., Heo, S., Jang, H., et al (2025). Machine learning-based early prediction of asthma in preschoolers: The COCOA birth cohort study. Pediatric Allergy and Immunology. https://doi.org/10.1111/pai.70223
29. Reis, J., Costa, R., Silva, F., et al (2025). Predicting Asthma Hospitalizations from Climate and Air Pollution Data: A Machine Learning-Based Approach. Climate. https://doi.org/10.3390/cli13020023
30. Custovic, D., Custovic, A., Gern, J., et al (2026). Model-based machine learning to understand the progression of asthma throughout the life course. Journal of Allergy and Clinical Immunology. https://doi.org/10.1016/j.jaci.2025.09.003
31. Liu, M., Cheng, Z., Sun, B. (2024). Utilizing deep machine learning approach to disclose hematological indices and biomarkers in asthma-COPD overlap. Clinica Chimica Acta. https://doi.org/10.1016/j.cca.2024.118796