



Artificial Neural Network-Based Feedback Control Strategy for Epidemiological SIR and SEIR Models

Oussama Chayoukh*¹, Omar Zakary¹, Mariam Redouane², Aadil Lahrouz²

¹ Ben M'Sick Faculty of Science, Hassan II University of Casablanca, Casablanca, Morocco.

² Faculty of Science and Techniques, Abdelmalek Essaadi University, Tangier, Morocco.

ABSTRACT: We investigate a branched artificial neural network (ANN) feedback controller for mitigating outbreaks in compartmental epidemic models. The architecture couples a shared trunk with two specialized branches and employs Soboleva–modified hyperbolic tangent (SMHT) activations to approximate the shape and boundedness of analytic control laws, yielding smooth, non–bang–bang signals suited to implementable interventions. The network is trained offline in supervised fashion on synthetic SIR trajectories labelled by a control that steers the infected population toward a low terminal target over a fixed horizon. On unseen SIR scenarios, the learned policy lowers peak prevalence and shortens outbreak duration relative to uncontrolled dynamics. When compared against simple baselines; however, the ANN achieves these outcomes with markedly smoother profiles and reduced actuation effort (time integral of the control), a property desirable for practice. Without retraining, the controller transfers to SEIR and retains qualitative benefits consistent with partial observability induced by the latent exposed class. We evaluate our suggested controller against conventional neural network baselines through ablation studies and robustness tests incorporating multiplicative process noise. The results demonstrate that our branched architecture reduces the attack size and peak infection with a comparable control effort. Importantly, the controller exhibits smooth, bounded actuation signals even when subjected to significant uncertainty. These results support learned feedback as a viable, effort-aware alternative to classical designs, with advantages expressed in the quality and efficiency of the intervention trajectory, a benefit beyond simply reducing the final number of cases. We discuss limitations and outline extensions: identification from data, observer design for latent/noisy states, explicit resource and rate constraints, and online adaptation under distribution shift.

Review History:

Received: Aug. 23, 2025

Revised: Nov. 04, 2025

Accepted: Dec. 11, 2025

Available Online: Dec. 31, 2025

Keywords:

Artificial Neural Network

Control Design

Epidemiological Models

Data-Driven Control

Branched Architecture

1- Introduction

One of the pillars of epidemiology is mathematical modeling, which offers essential understanding about the transmission, control and management of infectious diseases [1,2]. Such models allow researchers and policymakers to simulate different outbreak scenarios, predict epidemic peaks, and evaluate the potential impact of public health measures before they are implemented in real populations. Among these models, the traditional Susceptible-Infected-Recovered (SIR) framework, first established by the pioneers Kermack and McKendrick [3], is renowned for its simplicity and its ability to model disease dynamics by monitoring individuals' transitions between health states. Despite its simplicity, the SIR model has served as a foundational tool for both theoretical epidemiology and practical outbreak response, influencing decades of research in the field. However, for chronic diseases or those requiring consideration of demographic processes, including births and deaths, demographic elements must

be incorporated to accurately capture population dynamics over time [4–8]. Such considerations lead to demographic SIR models, widely explored yet still foundational for demonstrating the effectiveness of innovative control methods within a simpler framework before applying them to more complex scenarios. These demographic extensions enable more realistic simulations over long time horizons, making them particularly relevant for diseases that persist in populations for years or decades.

Recent studies, including our own prior work [9], have explored the possibility of transferring control strategies derived for simpler, non-demographic SIR models to more realistic demographic counterparts. We introduced a calibration framework based on the concept of τ -equivalence. With this, we showed that non-demographic models can accurately replicate the controlled trajectories of their demographic variants, particularly under moderate control intensities. This approach not only reduces the analytical and computational burden associated with directly optimizing complex demographic models, but also facilitates the reuse of

*Corresponding author's email: chayoukhossama@gmail.com



control strategies across structurally different epidemiological systems. The insights gained from such transferability studies directly inform the design of adaptable, data-driven controllers (such as the ANN-based framework presented here) capable of operating effectively in both simplified and demographically enriched contexts.

Controllability in epidemiology refers to the ability to guide an epidemiological system from an initial state to a desired health outcome through strategic interventions, such as isolation, quarantine, vaccination, and treatment [10]. This concept is closely linked to control theory in mathematics, where the focus is on determining whether—and how—specific interventions can steer a system toward a target trajectory. Typical objectives in disease control include reducing the size of the infected population over time, delaying the epidemic peak to avoid overwhelming healthcare systems, or flattening the infection curve within a specific timeframe [11,12]. Optimal control theory has been notably effective across multiple infectious diseases, including tuberculosis [13–17], malaria [18,19], human immunodeficiency virus (HIV) [20–28], hepatitis [29,30], and viral outbreaks such as influenza and COVID-19 [31–34]. In many of these cases, optimal control formulations have enabled the design of intervention schedules that account for both epidemiological impact and practical constraints such as resource availability and cost-effectiveness. Specifically, the Pontryagin Minimum Principle (PMP) [35] has been broadly used to identify optimal strategies balancing intervention costs and infection minimization [36–39]. PMP provides a mathematical framework for deriving necessary conditions for optimality and has been instrumental in shaping theoretical strategies for managing outbreaks under dynamic conditions.

Despite these successes, optimal control solutions derived through traditional methods are frequently computationally intensive, limiting their real-time applicability during rapidly evolving outbreaks [40]. This challenge becomes especially critical in the context of emerging diseases, where decisions must be made within days or even hours, and where parameter estimates may change as new data becomes available. Artificial intelligence, particularly artificial neural networks (ANNs), offers promising alternatives by significantly reducing computational demands and enhancing adaptability through their ability to learn and generalize optimal feedback controls directly from data [41–44]. Unlike classical optimal control approaches, ANNs can be trained offline and later deployed to produce near-instantaneous control outputs, making them particularly attractive for real-time epidemic management. Furthermore, their capacity to capture nonlinear relationships and adapt to varying epidemiological contexts enables them to remain effective even when model parameters shift over time.

In our earlier work [45], we directly compared PMP-based control to ANN-based feedback control for a nonlinear SIR system. That study highlighted PMP's sensitivity to local minima and computational demands, particularly when applied to nonlinear epidemic models, while demonstrating that ANNs could learn effective control policies capable

of overcoming some of these limitations. By training on simulated epidemic trajectories, the ANN-based approach achieved competitive or superior performance in infection suppression, with the added advantage of producing near-instantaneous control outputs once trained. These findings provided the initial motivation for the present work, in which we extend the ANN control concept to a branched architecture and explore its adaptability across different model structures, including SEIR dynamics.

In this paper, we present a specialized branched ANN control framework tailored to epidemiological models, specifically targeting the classical SIR and extended Susceptible-Exposed-Infected-Recovered (SEIR) models. Our network uses the pre-specified control policy to emulate our latest PMP-based neural network architecture, then is deployed to scenarios beyond those used for training. Our ANN is designed explicitly to process epidemiological states dynamically, generating control signals that effectively minimize infections and steer the system toward a predefined desired state at a specified final time. The branching design allows for the simultaneous handling of shared epidemiological features and model-specific characteristics, improving flexibility and adaptability.

We critically assess the ANN's performance through comparative analyses against uncontrolled scenarios, evaluating its practical potential and efficiency. These evaluations are conducted under diverse parameter configurations to highlight the robustness and generalization capacity of the learned control policies. Additionally, this study contributes to ongoing discourse regarding the strengths, limitations, and future possibilities of employing neural networks for epidemiological control. By positioning ANN-based control as both a complement and an alternative to traditional optimal control approaches, we aim to bridge the gap between theoretical models and actionable, data-driven outbreak management tools.

The rest of the paper is organized as follows. Section 2 (Methods) presents the model formulation, data generation, network architecture, training procedure, simulation setup, the baselines and ablation protocol, and the robustness protocol with multiplicative noise. Section 3 (Results and Interpretation) reports the training dynamics, SIR control performance, transfer to SEIR, the ablation against a simple MLP and a branched ReLU–sigmoid baseline, and the robustness analysis. Section 4 concludes and discusses limitations and extensions, including identification from data, observers for latent or noisy states, and explicit resource and rate constraints. Section 5 lists the nomenclature, and Section 6 provides the references.

2- Methods

2- 1- Epidemiological Model Formulation

The epidemic dynamics are captured by the Susceptible-Infected-Recovered (SIR) system shown in (1), where S represents the susceptible population, I the infected one, and R the recovered one. Here, $N = S + I + R$ represents the total population. The parameters β and γ represent the

infection and recovery rates, while $u(t)$ is an intervention (e.g., vaccination or isolation) designed to reduce infections. The main objective is to lower the infected population:

$$\begin{cases} \frac{dS}{dt} = -\frac{\beta SI}{N}, \\ \frac{dI}{dt} = \frac{\beta SI}{N} - (\gamma + u(t))I, \\ \frac{dR}{dt} = (\gamma + u(t))I. \end{cases} \quad (1)$$

This classic SIR model has been adapted to a wide range of real-world contexts, such as incorporating time delays and general incidence rates to better capture latency effects in disease progression [46], modelling demographic processes and infection-age structures for more accurate long-term forecasts [47], and integrating human mobility patterns to improve predictive accuracy for pandemics like COVID-19 [48]. These extensions have made the SIR framework relevant for analyzing and controlling diverse outbreaks, from local epidemics to global health crises.

For extended evaluation and to assess generalization, the Susceptible-Exposed-Infected-Recovered (SEIR) variant in (2) adds an Exposed compartment E , representing those infected but not yet infectious. Here, σ denotes the rate at which individuals transition from E to I :

$$\begin{cases} \frac{dS}{dt} = -\frac{\beta SI}{N}, \\ \frac{dE}{dt} = \frac{\beta SI}{N} - \sigma E, \\ \frac{dI}{dt} = \sigma E - \gamma I - u(t)I, \\ \frac{dR}{dt} = (\gamma + u(t))I. \end{cases} \quad (2)$$

SEIR models have also been successfully applied to a variety of epidemiological scenarios, including optimizing treatment strategies for tuberculosis in South Korea [17], integrating socio-psychological factors with genetic algorithms to predict COVID-19 resurgence [49], and quantifying the attenuation stage of infectious diseases to evaluate the effectiveness of regional control measures [50]. These applications demonstrate the adaptability of SEIR-type models in addressing both disease-specific dynamics and broader public health challenges.

2- 2- Data Generation

A dataset was constructed by integrating the SIR model over a grid of infection rates β (0.01 to 0.7, step 0.05) and recovery rates γ (0.01 to 0.5, step 0.05). Initial states (S, I, R) were systematically generated by incrementing S and I while keeping an eye on the proportions inside our

population of size N , then letting $R = N - (S + I)$. Each configuration was integrated in MATLAB using a fourth-order Runge-Kutta (RK4) solver over 40 days at 0.1-day intervals.

The control law $u(t)$, from [51], is used to steer $I(t)$ toward a desired level within a fixed horizon T . A parameter $\epsilon = 0.5$ modulates the sharpness of this intervention. As shown in (3),

$$u(t) = \frac{\beta S_0}{N} \left(\frac{a + e^{(\gamma+\epsilon)t}}{(1+a)e^{(\gamma+\epsilon)t}} \right)^{\beta I_0(1+a)/a(\gamma+\epsilon)N} - \gamma + \frac{(\gamma + \epsilon) e^{(\gamma+\epsilon)t}}{e^{(\gamma+\epsilon)t} + a} \quad (3)$$

where β represents the infection rate, γ denotes the recovery one, and A is a constant shaping the control function. Our objective is to train our designed neural network on a closed-form control law. This analytic solution provides a direct mapping from epidemiological states to intervention actions, simplifying data generation and ensuring consistent “optimal” labels for supervised learning. While precisely achieving I_d at T may be idealized, using this targeted control still ensures meaningful infection reduction and yields a well-defined reference for network training. By focusing on a standard formula, we integrate existing theoretical results into a data-driven framework, thereby retaining mathematical rigor while demonstrating practical generalization of the controller to diverse epidemic conditions.

Any generated control values of $u(t)$ falling outside the practical range $[0,1]$ were excluded from the dataset. The final dataset $[S, I, R, u]$ over all parameters and initial conditions was saved in a MATLAB matrix for subsequent neural network training.

2- 3- Neural Network Architecture

A branched neural network was designed to approximate the optimal control law based on the epidemic state. As illustrated in Fig. 1, the architecture consists of a trunk network T that processes the input vector $[S, I, R]$, two parallel branches $(B_i^{(j)})$ that specialize in distinct facets of the control dynamics, and a set of combining layers $(C^{(j)})$ that produce the final continuous control output. The trunk initially encodes the input through fully connected layers (500 neurons, Rectified Linear Unit (ReLU); 300 neurons, tanh) before splitting into two scalars that feed separate branches. Each branch then refines this representation, ultimately concatenating their outputs in a combining block culminating in a regression layer that yields a single real-valued control action.

In both branches, a Soboleva Modified Hyperbolic Tangent (SMHT) activation enables adaptive transformations that closely mimic the nonlinear structure of analytically derived control. The SMHT function is defined in (4):

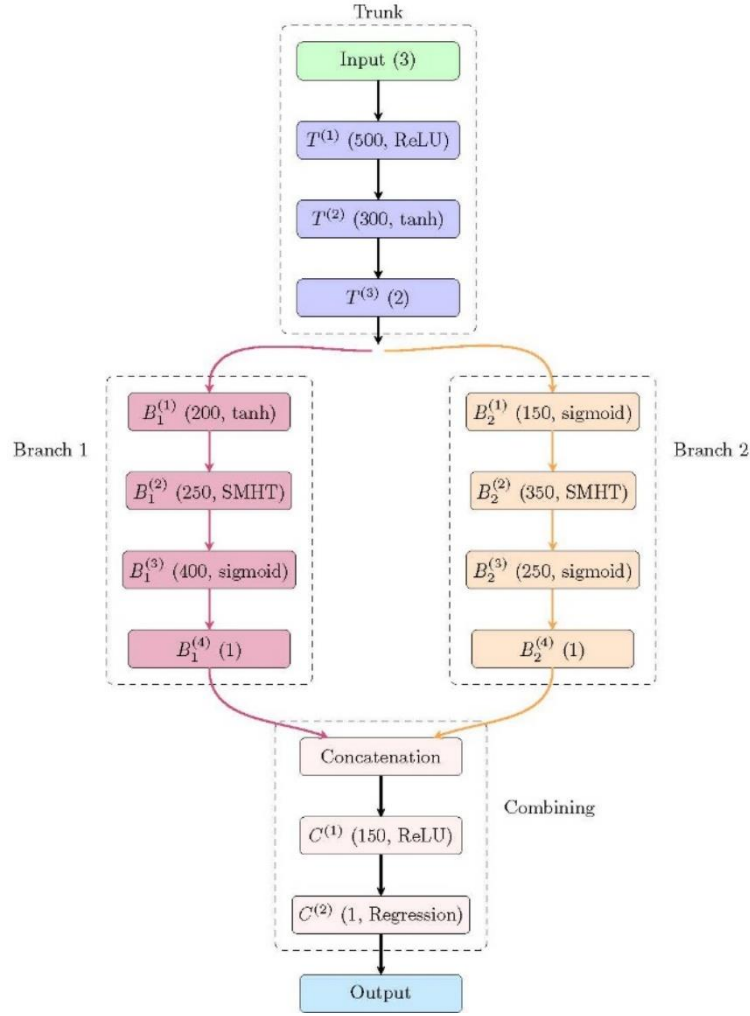


Fig. 1. Architecture of the branched neural network that includes trunk layers $T^{(i)}$, branch layers $B_i^{(j)}$, and combining layers $C^{(j)}$. Here, i is the branch number, and j is the layer index within each subgroup. Each layer is labeled as (“neurons”, “activation function”), with T for trunk, B_1 and B_2 for branches, and C for combining. The final regression layer has no nonlinear activation, producing continuous-valued outputs for regression tasks.

$$SMHT(x) = \frac{e^{ax} - e^{-bx}}{e^{cx} + e^{-dx}}, \quad (4)$$

where a and b introduce asymmetry, while c and d determine growth rate and saturation. Each branch uses different SMHT parameters to capture a broader range of potential epidemic scenarios, given that the underlying control requirements may shift with infection, recovery, or initial state variations. Although the branches share a similar layer structure, this parameter diversity allows them to specialize in complementary control regimes rather than duplicating

efforts. In particular, they ensure that abrupt or multi-regime control behaviors—such as those arising from changing ϵ or distinct target states—are effectively approximated over a wide range of epidemic conditions. Employing this branched architecture rather than a conventional multilayer perceptron enhances the capacity to capture complex, piecewise dynamics inherent in optimal interventions.

We adopted a capacity-ascending search to set the number of layers and neurons. Starting from a minimal baseline, we increased depth and width only when doing so improved validation performance and task metrics. The resulting architecture is thus the minimal-sufficient design meeting the objective, with restrained complexity.

Table 1. Setup conditions and parameters used in the six configurations. Configurations 1–3 correspond to the SIR model, while configurations 4–6 extend to the SEIR model.

Config	S_0	I_0	E_0	R_0	β	γ	σ
1	86	14	–	0	0.40	0.01	–
2	45	35	–	20	0.51	0.02	–
3	52	38	–	10	0.37	0.01	–
4	74	10	16	0	0.40	0.01	0.30
5	37	33	25	5	0.51	0.05	0.10
6	52	19	26	3	0.37	0.20	0.20

2- 4- Training Procedure

The dataset was randomly partitioned into 70% for training, 15% for validation, and 15% for testing, ensuring class balance across subsets. Normalization to zero mean and unit variance was applied feature-wise using (5) and the training set statistics only:

$$x_{\text{norm}} = \frac{x - \mu_{\text{train}}}{\sigma_{\text{train}}}, \quad (5)$$

where μ_{train} and σ_{train} denote the mean and standard deviation of each feature in the training subset. This mitigated scale discrepancies and improved numerical stability.

The network underwent supervised learning to minimize the mean squared error (MSE) between predicted controls and their reference values. MSE was computed per mini-batch and averaged over all time steps. We employed the Adam optimizer which commenced with a learning rate of 0.001, which decayed by a factor of 0.1 after every 10 epochs. Training progressed for a maximum of 30 epochs and was set to be halted prematurely by early stopping if the validation performance did not improve for 10 consecutive epochs.

The 30 epochs were partitioned with mini-batches of size 1024, balancing convergence stability and computational efficiency. An L_2 regularization term ($\lambda = 0.01$) was added to the loss on the combining layers to reduce overfitting.

2- 5- Simulation

All simulations were conducted using MATLAB with a fourth-order Runge–Kutta (RK4) solver, maintaining consistency with the data generation methodology. The integration time step was set to 0.1 days over a 40-day simulation horizon, providing sufficient temporal resolution to capture rapid epidemic dynamics while maintaining computational efficiency.

To facilitate numerical stability and reduce computational complexity, we normalized the population to $N = 100$, representing 100% of the population in percentage terms. This normalization does not affect the relative dynamics

of the system while simplifying the computational burden during both training and evaluation phases.

Six distinct epidemic scenarios were designed to comprehensively evaluate the ANN controller's performance across varying epidemic conditions and model structures. These configurations were strategically selected to span different epidemic regimes:

- Configurations 1–3: SIR model scenarios with varying transmission and recovery dynamics.
- Configurations 4–6: SEIR model scenarios to test cross-model generalization capabilities.

The configurations encompass a range of epidemic severity levels, from moderate transmission scenarios to high-risk situations with potential for explosive growth. Table 1 presents the specific initial conditions and epidemiological parameters for each configuration.

The trained neural network controller was implemented using a closed-loop feedback strategy, where the control action $u(t)$ was computed dynamically at each temporal increment based on the instantaneous epidemic state. This approach simulates realistic real-time intervention scenarios where control decisions must be made based on observed conditions.

2- 6- Baselines & Ablation Protocol

To evaluate our results, we compared our network to two baselines whose architectures are shown in Fig. 2 (simple non-branched MLP, ReLU/Sigmoid) and Fig. 3 (branched ReLU/Sigmoid) respectively, to conduct an ablation study. Layer widths and depths were fixed after a small sweep to identify the smallest configurations that met our target accuracy, so the ablation isolates architectural effects rather than raw capacity. For fairness, both models use the same training/validation split and preprocessing, identical epochs, optimizer with the same learning-rate schedule, mini-batch size, etc.

The comparison of the networks was based on metrics such as: Peak infection I_{max} , time-to-peak t_{peak} , attack size $\int_0^{t_f} I(t) dt$, and control energy $\int_0^{t_f} u(t)^2 dt$.




Fig. 2. Architecture of the simple neural network that includes activation functions such as ReLU and sigmoid function.

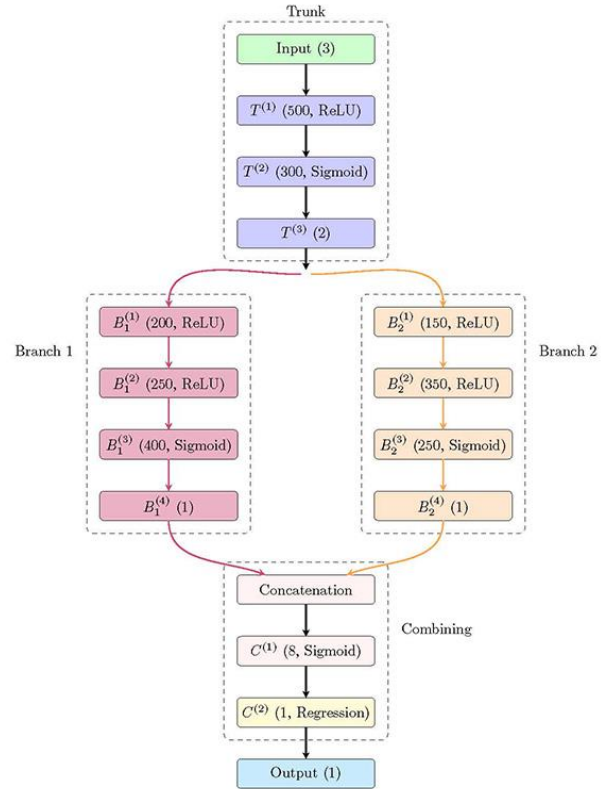


Fig. 3. Architecture of the branched neural network that includes usual activation functions.

2- 7- Robustness Protocol

We tested the controller under stepwise multiplicative noise on the model rates to examine its robustness and capacity to sustain acceptable performance when the model and data are subject to fluctuations and noise. At each Runge-Kutta step k , every perturbed rate p is updated as

$$p_k = p \exp(\eta \xi_k), \quad \xi_k \sim \mathcal{N}(0,1),$$

where p is the nominal value. We used one noise scale η shared by all perturbed rates. For SIR we perturbed the transmission rate β and the recovery rate γ . For SEIR we perturbed a , r , and the incubation rate σ . We injected no measurement noise: the policy sees the true state. We tested three levels, $\eta \in \{0, 0.05, 0.10\}$. For each model we ran the same set of initial conditions and parameter cases across several random seeds and report mean \pm standard deviation for the attack size, the control energy, the infection peak and its time, and the mass error $\max_k |S_k + E_k + I_k + R_k - N|$ to check population conservation; in SIR the E_k term is omitted. We reported an admissibility rate given by the proportion of runs with $u(t) \in [0,1]$ for all t .

3- Results and Interpretation

3- 1- Training Dynamics

Fig. 4 reports the training and validation loss over 30 epochs. The curves show a sharp descent during the first few epochs, followed by a slower, approximately exponential regime and a clear plateau in the last third of training. A small transient bump in the validation trace around the mid-epochs is quickly damped, after which both traces settle to a near-stationary level.

The two curves remain closely aligned throughout, with only a negligible generalization gap. This behavior indicates that the model capacity and regularization are well matched to the task: the network fits the training data without memorization and retains predictive performance on held-out examples as best as it can. The mild mid-training oscillations are consistent with stochastic optimization noise rather than the onset of overfitting, given the subsequent return to the prevailing downward trend and the final plateau.

3- 2- Neural Network Control Performance for the SIR Model

The trained neural network controller demonstrated strong performance in reducing infection rates across diverse epidemic scenarios.




Fig. 4. Training and validation loss curves for the suggested ANN model.

Fig. 5 presents results for the SIR model dynamics under the configurations 1-3. In configuration (a), the uncontrolled scenario shows a dramatic increase in infections, peaking at approximately 85 cases before slowly declining. With ANN control, this peak was completely averted, keeping infections consistently below 10. Similar performance was observed in configurations (b) and (c).

The corresponding population dynamics for these high-risk scenarios (Fig. 6) reveal an important characteristic of the neural network controller’s strategy. In all three configurations, the controller maintained a significantly higher susceptible population compared to uncontrolled scenarios. This suggests that rather than simply accelerating the transition through the infectious state, the controller effectively suppressed transmission while simultaneously promoting recovery, a more efficient approach to epidemic management.

The control signals for these scenarios (Fig. 7) demonstrate adaptive control approaches based on epidemic severity. Configuration (a) shows an initially high control value that drops sharply around day 7 before gradually increasing again, suggesting a pulsed intervention strategy. Configurations (b) and (c) exhibit more conventional monotonically increasing control signals, although with different rates of increase reflecting the specific dynamics of each scenario.

3- 3- Generalization to SEIR Model

A critical test of the neural network controller’s robustness was its application to the SEIR model, despite being trained exclusively on SIR dynamics. Fig. 8 shows the performance on configurations 4-6, which introduced an exposed compartment that was not present during training. The controller continued to outperform uncontrolled scenarios,

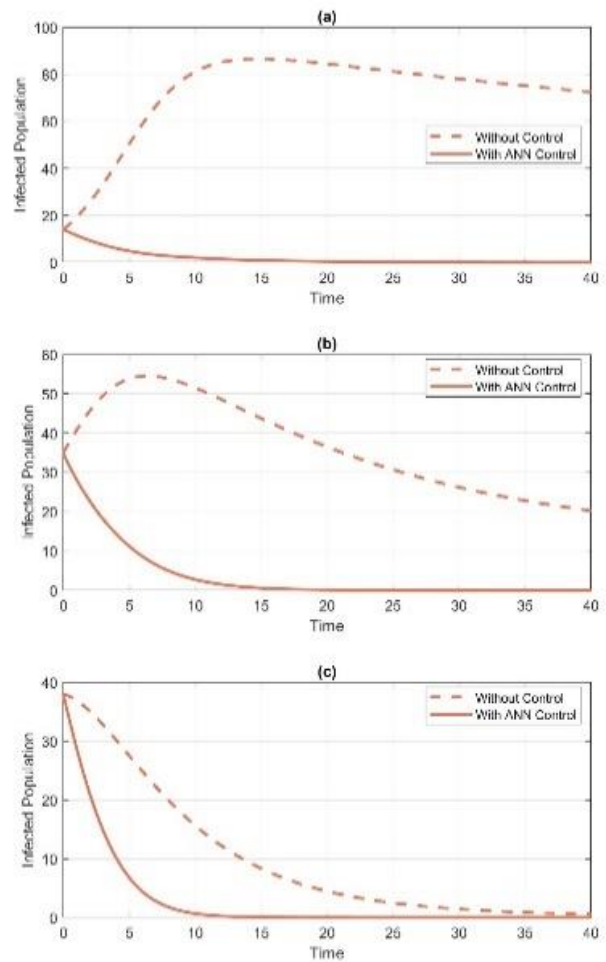


Fig. 5. Comparison of infected population dynamics in high-risk epidemic configurations: (a) Configuration 1; (b) Configuration 2; (c) Configuration 3.




Fig. 6. Susceptible and recovered population dynamics. (a) Configuration 1; (b) Configuration 2; (c) Configuration 3.

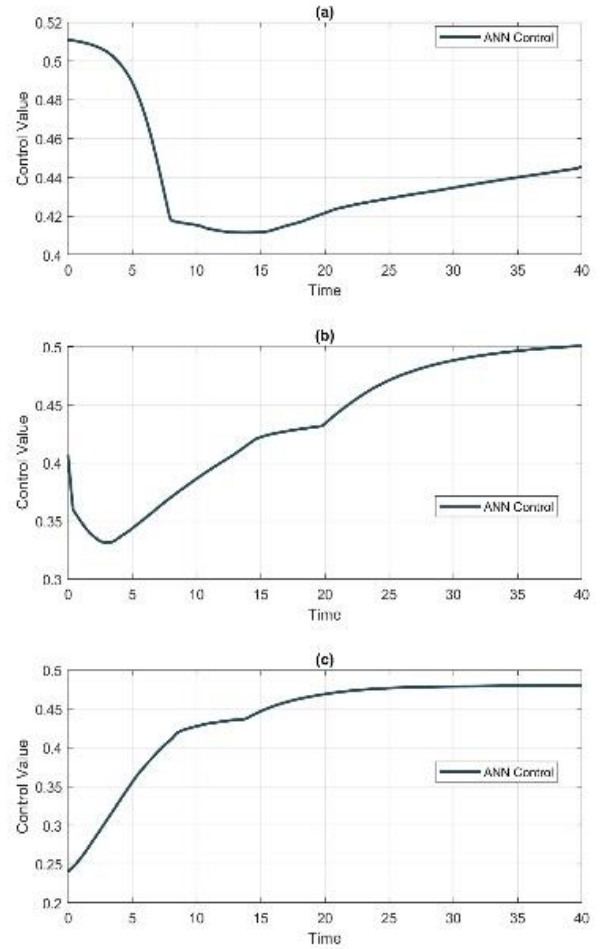


Fig. 7. Control signals generated by the ANN for the high-risk configurations 1-3.

although with less dramatic reductions in peak infections compared to the SIR cases. This indicates successful transfer learning across model structures, an important capability for practical implementation where the true epidemic model may be uncertain.

Fig. 9 presents the remaining compartmental dynamics for the SEIR model, including the exposed population not visible to the controller. Despite this limitation, the controller effectively managed all compartments. One notable observation is the controller’s indirect influence on the exposed compartment due to the nature of dynamical systems, which was consistently reduced in controlled scenarios compared to uncontrolled ones.

As in SIR, control profiles are mostly monotonically increasing and smooth; SEIR (Fig. 10) adds mild adjustments around days 15–20 in (a) and (c), suggesting a more conservative response to latent-period uncertainty. Possibly, this indicates a more conservative approach when dealing with the additional uncertainty introduced by the unobserved exposed compartment.

Taken together, the SIR and SEIR experiments indicate that the learned controller operates as a smooth, bounded, intensity-modulated feedback that reliably reduces prevalence while preserving a larger susceptible pool. Its effectiveness carries over to SEIR despite training solely on SIR, consistent with model-agnostic policies that depend only on observable (S, I, R) ; however, the gains are smaller when the latent period dominates, reflecting limits from partial observability. Practically, the graded (non-bang-bang) profiles map naturally to proportional vaccination or contact-reduction schedules, and the bounded output induced by the SMHT-based architecture is compatible with effort-aware implementations that avoid abrupt actuation. These properties align with established insights in epidemic control and implementable feedback design.

3-4 Ablation Results

Fig. 11 and Fig. 12 illustrate the training and validation loss over 30 epochs for the baseline models. The curves follow a pattern similar to the learning curves in Fig. 4,

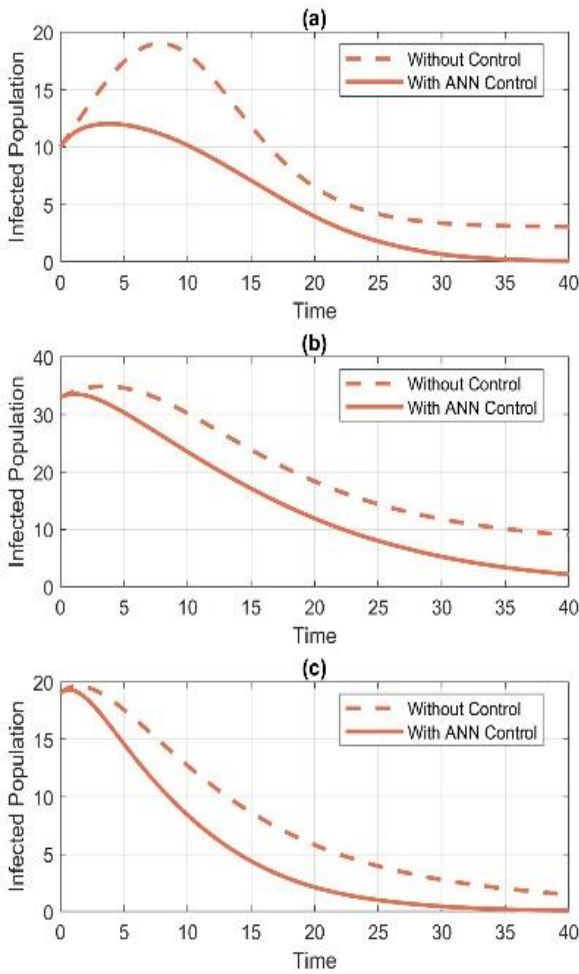


Fig. 8. Infected population dynamics for SEIR model configurations: (a) Configuration 4; (b) Configuration 5; (c) Configuration 6.

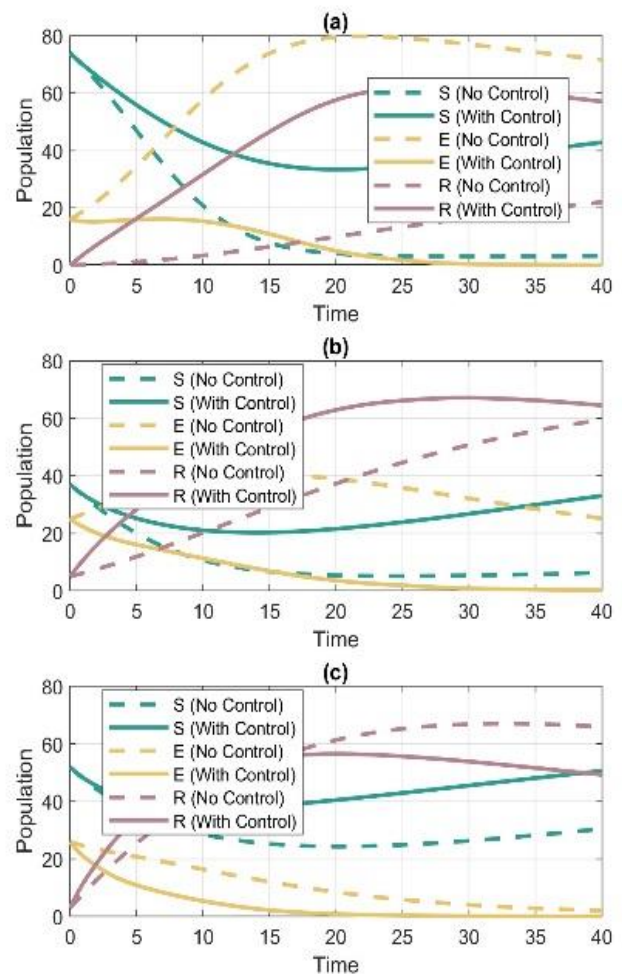


Fig. 9. The remaining compartmental dynamics for SEIR model configurations 4-6, showing susceptible (S), exposed (E), and recovered (R) populations with and without ANN control.

starting with a fast decrease and then moving toward a stable region. This behavior suggests that the baseline networks have converged properly and that the training process was effective. The similarity in shape and stability with the reference curves indicates that the models have learned the underlying dynamics in a consistent and reliable way.

Here we ablate architectural choices by comparing the proposed branched network with SMHT activations to a branched network with standard activations and to a simple MLP, under the same experimental setups (in Table 1).

Table 2 presents the peak infection for all configurations. As expected, the SIR cases peak at their initial condition because the control acts immediately and the SIR structure has no latent stage. The separation between models appears in the SEIR cases. The proposed network attains the smallest peaks in all three SEIR configurations. The gains are modest in absolute terms, typically between 0.1% and 1.0% points of the infected population, but they are consistent across cases. On average, peak infection is 25.12% for the proposed

network versus 25.35% for the branched baseline and 25.38% for the simple network, which is about a one percent reduction relative to the comparators.

Time to peak complements this picture. Table 3 shows that the SIR model peak at time zero for all configurations. In the SEIR configurations, the proposed network reaches the peak earlier and with a smaller size. The mean time to peak is 0.48 for the proposed model compared to 1.00 for the branched baseline and 1.12 for the simple network. This is a reduction of about 52% and 57%, respectively, which shortens the period of high prevalence and accelerates the decline phase.

Cumulative burden is captured by the attack size. Table 4 reports a clear and robust advantage for the proposed network. The mean attack size over all six configurations is 198.51 for the proposed model, compared to 246.74 for the branched baseline and 252.90 for the simple network. This corresponds to reductions of about 19.5 percent and 21.5 percent, respectively. The only exception is configuration 1, where the attack size of the proposed model is higher than

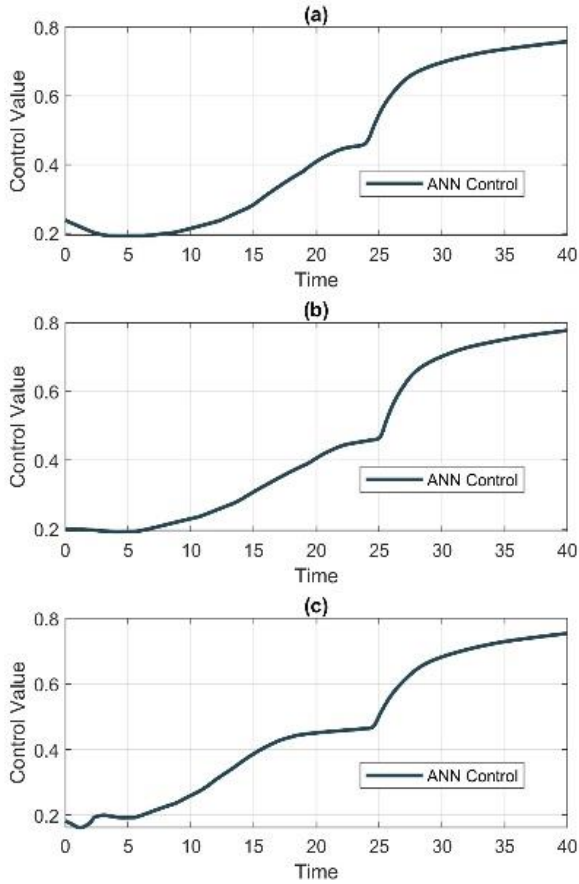


Fig. 10. Control signals generated by the ANN for SEIR model configurations 4-6.

both baselines. Outside this outlier, the proposed architecture yields consistent reductions across SIR and SEIR cases.

Energy expenditure of the control is summarized in Table 5. The proposed model uses slightly more energy than the branched baseline on average, 7.65 versus 7.65, which is a difference of about 0.01 in absolute terms and roughly 0.1 percent. Compared to the simple network, the proposed model requires slightly more energy, 7.65 versus 6.88 on average, which is about 11% higher. Taken together with Table 4, these results show that a negligible increase relative to the branched baseline in control energy is associated with a large drop in attack size, and that the gains in attack size and time to peak also hold against the simple network despite its lower energy cost.

In summary, across matched settings, the proposed branched SMHT architecture produces the smallest peaks in SEIR, reaches those peaks earlier, and reduces the cumulative attack size by about one fifth on average, with only a minimal increase in energy relative to the branched baseline. The simple network expends less energy but performs worse on

peak size, time to peak, and attack size, which is consistent with its lower representational capacity.

The branched-SMHT design outperforms the baselines because the SMHT activation is smooth, bounded, and slope-tunable, preserving gradient quality at low amplitudes and preventing the saturation and dead-zone behavior seen with hard nonlinearities. This lets the network represent “soft switching” control laws that rise quickly but remain stable, which matches the needs of epidemic mitigation where early, moderate action limits later surges (consistent with the smaller peaks and earlier peaks in Tables 2–3). The branched architecture further separates shared epidemiological features (learned in the trunk) from regime-specific responses (handled in the branches). One branch can specialize in fast transient shaping (anticipating the exposed-to-infected spillover in SEIR) while another emphasizes longer-horizon containment and energy efficiency. This division reduces interference between competing objectives and acts like an implicit gain-scheduling policy across configurations, which explains the lower attack size at nearly unchanged energy relative to the branched baseline. These properties align with prior observations that smooth, non-monotone activations enhance trainability and expressivity, and that multi-branch designs improve specialization without sacrificing shared structure

3- 5- Robustness Analysis

We tested the controller under multiplicative noise on the transmission rate. At each step the policy read the current state and produced a control action. We varied the noise scale $\eta \in \{0, 0.05, 0.1\}$. For every η we ran 20 random seeds per scenario and report mean \pm standard deviation across seeds and cases. The metrics are attack size, control energy $\int u^2 dt$, peak infected and its timing, mass error, and admissibility. The summary appears in Table 6.

We tested the controller under stepwise multiplicative noise in the epidemic dynamics. The same network, trained on clean data, was evaluated on SIR and SEIR settings. For each model we ran the same cases and seeds at three noise levels $\eta = 0, 0.05$, and 0.1 . At every level we measured the attack size, the control energy, the infection peak and its time, and the mass error to check population conservation. We also recorded the admissibility rate, that is the percentage of runs with $u(t)$ inside $[0, 1]$ without any clipping. Results are summarized in Table 6 for SIR and Table 7 for SEIR.

On the SIR model the means of all metrics remain essentially unchanged when η increases from 0 to 0.1. The attack size changes from 1439.41 to 1439.39 with a standard deviation near 981. The energy stays near 6.56 with a standard deviation near 0.35. The infection peak is about 56.24 with a standard deviation near 19.80. The time to peak decreases slightly from about 9.77 to about 9.76, which is very small compared with its standard deviation near 9.07. The mass error is numerically zero for all settings. Admissibility is 100 percent for all runs, which shows that the controller outputs remain within $[0, 1]$ without any explicit bounding.

On the SEIR model we observe the same pattern. The attack size moves from 1409.75 to 1409.74 with a standard




Fig. 11. Training and validation loss curves for the simple ANN model.




Fig. 12. Training and validation loss curves for the branched ANN model with standard activation functions.

Table 2. Comparison of Peak Infection across networks (mean over all cases).

Model	Configuration	Branched SMHT	Branched	Simple
SIR	1	14	14	14
SIR	2	35	35	35
SIR	3	38	38	38
SEIR	4	11.0760	12.1010	12.2070
SEIR	5	33.4080	33.6250	33.6840
SEIR	6	19.2520	19.3450	19.3730
Mean		25.1227	25.3452	25.3773

Table 3. Comparison of time-to-peak t_{peak} across networks (mean over all cases).

Model	Configuration	Branched SMHT	Branched	Simple
SIR	1	0.0	0.0	0.0
SIR	2	0.0	0.0	0.0
SIR	3	0.0	0.0	0.0
SEIR	4	1.6	4.0	4.5
SEIR	5	0.8	1.3	1.4
SEIR	6	0.5	0.7	0.8
Mean		0.4833	1.0	1.1167

Table 4. Comparison of Attack Size $\int_0^{t_f} I(t) dt$ across networks (mean over all cases).

Model	Configuration	Branched SMHT	Branched	Simple
SIR	1	72.1090	56.1770	67.5300
SIR	2	145.3000	254.0100	260.0100
SIR	3	111.7000	155.5000	163.0700
SEIR	4	149.9800	210.0400	212.7500
SEIR	5	529.7300	596.0800	602.1400
SEIR	6	182.2700	208.6100	211.9200
Mean		198.5148	246.7362	252.9033

Table 5. Comparison of Energy $\int_0^{t_f} u(t)^2 dt$ across networks (mean over all cases).

Model	Configuration	Branched SMHT	Branched	Simple
SIR	1	14	14	14
SIR	2	35	35	35
SIR	3	38	38	38
SEIR	4	11.0760	12.1010	12.2070
SEIR	5	33.4080	33.6250	33.6840
SEIR	6	19.2520	19.3450	19.3730
Mean		25.1227	25.3452	25.3773

Table 6. Robustness of the controller on the SIR model under noise across cases and seeds.

η	Attack size (mean±sd)	Energy (mean±sd)	Peak Infection (mean±sd)	Time-to-peak (mean±sd)	Mass error (mean±sd)	Admissibility %
0	1439.41144±980.90321	6.56191±0.35132	56.24504±19.80422	9.77123±9.07015	0.00±0.00	100%
0.05	1439.40322±980.90211	6.56120±0.35682	56.24322±19.80003	9.76656±9.06766	0.00±0.00	100%
0.1	1439.39164±980.89001	6.56064±0.352201	56.24210±19.80106	9.76155±9.05543	0.00±0.00	100%

Table 7. Robustness of the controller on the SEIR model under noise across cases and seeds.

η	Attack size (mean±sd)	Energy (mean±sd)	Peak Infection (mean±sd)	Time-to-peak (mean±sd)	Mass error (mean±sd)	Admissibility %
0	1409.75036±871.62223	6.54267±0.29842	47.3276±25.06002	14.43641±10.49113	0.00±0.00	100%
0.05	1409.74923±871.61024	6.54176±0.29311	47.32134±25.05950	14.43550±10.49022	0.00±0.00	100%
0.1	1409.74134±871.61654	6.54079±0.29001	47.32001±25.05367	14.43459±10.48931	0.00±0.00	100%

deviation near 872. The energy is about 6.54 with a standard deviation near 0.30. The peak remains around 47.32 with a standard deviation near 25.06. The time to peak shifts from about 14.43 to about 14.43, again well within its standard deviation near 10.49. The mass error is zero across all settings. Admissibility is 100 percent for all runs.

Taken together, these results indicate that the learned controller is robust to multiplicative noise up to $\eta=0.1$. Performance does not degrade in any meaningful way and the small changes that appear are far smaller than the variability across cases captured by the reported standard deviations. Population conservation is preserved and there is no loss of control feasibility. This supports the reliability of the controller when the underlying rates fluctuate within the tested range.

4- Conclusion and Future Directions

This work has demonstrated that a tailored branched neural network, leveraging Soboleva-modified hyperbolic tangent activations and a modular trunk–branch–combine design, can emulate a precomputed analytic control law and apply it in real time to both SIR and SEIR frameworks. Across a diverse set of epidemic scenarios, the ANN consistently lowered peak prevalence, shortened epidemic duration, and preserved a larger susceptible pool, outperforming uncontrolled baselines and even generalizing to SEIR dynamics without retraining. An ablation against a simple non-branched MLP and a branched ReLU/sigmoid baseline supports this choice, with consistent gains in attack size, peak, and time-to-peak at comparable energy. A robustness study with stepwise multiplicative noise on transmission (η up to 0.10) showed essentially unchanged performance and 100% admissibility without clipping, indicating resilience to moderate parametric variability.

Nevertheless, our study remains bounded by several important limitations. First, the size and diversity of our training dataset (though sufficient to validate reduction of infections) fall short of fully capturing the manifold behaviors of optimal controls under highly heterogeneous parameter regimes; our parameter ranges were compact and T and I_d were held fixed for computational reasons, and scaling up to larger, more varied data would bolster generalization at the expense of greater computational demands. Second, by training solely on instantaneous state–control pairs and evaluating performance against a single outcome (“reduction of infected population”), we risk the network learning only a coarse proxy rather than the true functional form of the analytic control; incorporating richer information such as explicit time indices, horizon length T , and desired terminal targets I_d would deepen the network’s internalization of control dynamics and ensure it reproduces the subtleties

of the underlying feedback law. Third, our robustness tests perturb only the dynamics (no observation noise); extending to state-measurement noise and drift would strengthen external validity. Finally, our benchmarks all admit a natural disease-free equilibrium, meaning that any admissible continuous controller will eventually drive infections toward zero; a more stringent test would involve models without a stable disease-free state, forcing the ANN to learn genuinely nontrivial intervention schedules rather than relying on trivial decay.

Nomenclature

A	Constant Shaping the Control Law
E	Exposed Population
I	Infected Population
\mathcal{N}	Standard Normal Distribution
N	Total Population
R	Recovered Population
S	Susceptible Population
t	Time
T	Final Time Horizon
u	Control Function

Greek symbols

β	Infection Rate
ϵ	Control Sharpness Parameter
η	Noise Intensity Scale for Multiplicative Perturbations
γ	Recovery Rate
λ	Regularization Coefficient
σ	Transition Rate from Exposed to Infected
τ	Time-shift Parameter
ξ	Standard Normal Innovation in Noise Model

Subscript

o	Initial Value
d	Desired State
i	Branch Index in the Artificial Neural Network
$train$	Training Dataset Statistics

Superscript

$()$	Layer Index in the Artificial Neural Network
------	--

References

- [1] Brauer F, Castillo-Chavez C. *Mathematical Models in Population Biology and Epidemiology*. Springer, New York; 2012. (Texts Appl Math; 40). doi:10.1007/978-1-4614-1686-9.
- [2] Brauer F, van den Driessche P, Wu J, Morel J-M, Takens F, Teissier B, eds. *Mathematical Epidemiology*. Springer, Berlin; 2008. (Lecture Notes Math; 1945). doi:10.1007/978-3-540-78911-6.
- [3] Kermack WO, McKendrick AG. A contribution to the mathematical theory of epidemics. *Proc R Soc Lond A*. 1927;115(772):700–21. doi:10.1098/rspa.1927.0118.
- [4] Busenberg SN, Haderl KP. Demography and epidemics. *Math Biosci*. 1990;101(1):63–74.
- [5] Busenberg S, van den Driessche P. Analysis of a disease transmission model in a population with varying size. *J Math Biol*. 1990;28(3):257–70. doi:10.1007/BF00178776.
- [6] Derrick WR, van den Driessche P. A disease transmission model in a nonconstant population. *J Math Biol*. 1993;31(5):495–512. doi:10.1007/BF00173889.
- [7] Martcheva M, Castillo-Chavez C. Diseases with chronic stage in a population with varying size. *Math Biosci*. 2003;182(1):1–25.
- [8] Anderson RM, May RM. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, Oxford; 1991.
- [9] Chayoukh O, Zakary O. Transferring control strategies in epidemiological models using τ -equivalences. *Commun Math Biol Neurosci*. 2025;2025:54.
- [10] Sharomi O, Malik T. Optimal control in epidemiology. *Ann Oper Res*. 2017;251(1):55–71. doi:10.1007/s10479-015-1834-4.
- [11] Klamka J. Controllability of dynamical systems: a survey. *Bull Pol Acad Sci Tech Sci*. 2013;61(2):335–42.
- [12] Sussmann HJ, Jurdjevic V. Controllability of nonlinear systems. *J Differ Equ*. 1972;12(1):95–116.
- [13] Jung E, Lenhart S, Feng Z. Optimal control of treatments in a two-strain tuberculosis model. *Discrete Contin Dyn Syst B*. 2002;2(4):473–82. doi:10.3934/dcdsb.2002.2.473.
- [14] Moualeu DP, Weiser M, Ehrig R, Deuffhard P. Optimal control for a tuberculosis model with undetected cases in Cameroon. *Commun Nonlinear Sci Numer Simul*. 2015;20(3):986–1003.
- [15] Silva CJ, Torres DFM. Optimal control for a tuberculosis model with reinfection and post-exposure interventions. *Math Biosci*. 2013;244(2):154–64.
- [16] Augusto FB, Adekunle AI. Optimal control of a two-strain tuberculosis–HIV/AIDS co-infection model. *Biosystems*. 2014;119:20–44.
- [17] Whang S, Choi S, Jung E. A dynamic model for tuberculosis transmission and optimal treatment strategies in South Korea. *J Theor Biol*. 2011;279(1):120–31.
- [18] Kim BN, Nah K, Chu C, Ryu SU, Kang YH, Kim Y. Optimal control strategy of Plasmodium vivax malaria transmission in Korea. *Osong Public Health Res Perspect*. 2012;3(3):128.
- [19] Prosper O, Ruktanonchai N, Martcheva M. Optimal vaccination and bednet maintenance for the control of malaria in a region with naturally acquired immunity. *J Theor Biol*. 2014;353:142–56.
- [20] Joshi HR. Optimal control of an HIV immunology model. *Optim Control Appl Methods*. 2002;23(4):199–213. doi:10.1002/oca.710.
- [21] Fister KR, Lenhart S, McNally JS. Optimizing chemotherapy in an HIV model. *Electron J Differ Equ*. 1998.
- [22] Yang Y, Xiao Y, Wu J. Pulse HIV vaccination: feasibility for virus eradication and optimal vaccination schedule. *Bull Math Biol*. 2013;75(5):725–51.
- [23] Kwon HD, Lee J, Yang SD. Optimal control of an age-structured model of HIV infection. *Appl Math Comput*. 2012;219(5):2766–79.
- [24] Roshanfekar M, Farahi MH, Rahbarian R. A different approach of optimal control on an HIV immunology model. *Ain Shams Eng J*. 2014;5(1):213–19.
- [25] Zhou Y, Liang Y, Wu J. An optimal strategy for HIV multitherapy. *J Comput Appl Math*. 2014;263:326–37.
- [26] Adams BM, Banks HT, Davidian M, Kwon HD, Tran HT, Wynne SN, et al. HIV dynamics: modeling, data analysis, and optimal treatment protocols. *J Comput Appl Math*. 2005;184(1):10–49. doi:10.1016/j.cam.2005.02.002.
- [27] Costanza V, Rivadeneira PS, Biafore FL, D’Attellis CE. Optimizing thymic recovery in HIV patients through multidrug therapies. *Biomed Signal Process Control*. 2013;8(1):90–97.
- [28] Zakary O, Rachik M, Elmouki I. On the impact of awareness programs in HIV/AIDS prevention: an SIR model with optimal control. *Int J Comput Appl*. 2016;133(9):1–6.
- [29] Rong L, Perelson AS. Treatment of hepatitis C virus infection with interferon and small-molecule direct antivirals: viral kinetics and modeling. *Crit Rev Immunol*. 2010;30(2):131–48.
- [30] Zakary O, Rachik M, Elmouki I. On effectiveness of an optimal antiviral bitherapy in HBV-HDV coinfection model. *Int J Comput Appl*. 2015;127(12):1–10.
- [31] Augusto FB. Optimal isolation control strategies and cost-effectiveness analysis of a two-strain avian influenza model. *Biosystems*. 2013;113(3):155–64.
- [32] Yan X, Zou Y. Optimal and sub-optimal quarantine and isolation control in SARS epidemics. *Math Comput Model*. 2007;47(1–2):235–45.
- [33] Li Y. Optimal control for an epidemic model of COVID-19 with time-varying parameters. *Mathematics*.

- 2024;12(10):1484.
- [34] Seddighi Chaharborj S, Seddighi Chaharborj S, Hassanzadeh Asl J, Phang PS. Controlling of pandemic COVID-19 using optimal control theory. *Results Phys.* 2021;26:104311.
- [35] Pontryagin LS. *Mathematical Theory of Optimal Processes.* CRC Press, Boca Raton; 1987.
- [36] Bolzoni L, Bonacini E, Soresina C, Groppi M. Time-optimal control strategies in SIR epidemic models. *Math Biosci.* 2017;292:86–96.
- [37] Grigorieva EV, Khailov EN, Korobeinikov A. Optimal control for a SIR epidemic model with nonlinear incidence rate. *Math Model Nat Phenom.* 2016;11(4):89–104.
- [38] Bakare EA, Nwagwo A, Danso-Addo E. Optimal control analysis of a SIR epidemic model with constant recruitment. *Int J Appl Math Res.* 2014;3(3):273–85.
- [39] Zakary O, Rachik M, Elmouki I. On the analysis of a multi-regions discrete SIR epidemic model: an optimal control approach. *Int J Dyn Control.* 2017;5(3):917–30. doi:10.1007/s40435-016-0233-2.
- [40] Lenhart S, Workman JT. *Optimal Control Applied to Biological Models.* Chapman & Hall/CRC, Boca Raton; 2007.
- [41] Goodfellow I, Bengio Y, Courville A. *Deep Learning.* MIT Press, Cambridge, MA; 2016.
- [42] Sutton RS, Barto AG. *Reinforcement Learning: An Introduction.* MIT Press, Cambridge, MA; 2018.
- [43] LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature.* 2015;521(7553):436–44.
- [44] Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, et al. Mastering the game of Go with deep neural networks and tree search. *Nature.* 2016;529(7587):484–89.
- [45] Chayoukh O, Zakary O. Application of artificial intelligence to control a nonlinear SIR model. In: Chakir A, Andry JF, Ullah A, Bansal R, Ghazouani M, eds. *Engineering Applications of Artificial Intelligence.* Springer, Cham; 2024. p. 23–39. doi:10.1007/978-3-031-50300-9_2.
- [46] Hattaf K, Lashari A, Louartassi Y, Yousfi N. A delayed SIR epidemic model with a general incidence rate. *Electron J Qual Theory Differ Equ.* 2013;2013(3):1–9.
- [47] Jing W, Jin Z, Zhang J. An SIR pairwise epidemic model with infection age and demography. *J Biol Dyn.* 2018;12(1):486–508.
- [48] AlQadi H, Bani-Yaghoub M. Incorporating global dynamics to improve the accuracy of disease models: example of a COVID-19 SIR model. *PLoS One.* 2022;17(4):e0265815.
- [49] Wang H, Wu D, Luo J, Zhang J. Integrating socio-psychological factors in the SEIR model optimized by a genetic algorithm for COVID-19 trend analysis. *Sci Rep.* 2024;14:15684.
- [50] Sun T, Jin B, Wu Y, Bao J. A study of the attenuation stage of a global infectious disease. *Front Public Health.* 2024;12:1379481.
- [51] Zakary O, Bidah S, Rachik M. Optimizing infection trajectories: innovation in controllability of nonlinear SIR model. *Rev Mex Ing Biomed.* 2024;45(2):151–71.

HOW TO CITE THIS ARTICLE

O. Chayoukh, O. Zakary, M. Redouane, A. Lahrouz, *Artificial Neural Network-Based Feedback Control Strategy for Epidemiological SIR and SEIR Models*, *AUT J. Model. Simul.*, 57(2) (2025) 157-172.

DOI: [10.22060/miscj.2026.24611.5428](https://doi.org/10.22060/miscj.2026.24611.5428)

