Minimax Theory and its Applications

Volume 11 (2026), No. 2s, 301 to 320

A Minimax Differential Game Framework for Personalized Closed-Loop Therapeutic Control in Smart Healthcare Systems

A.Sabah Afroze[1], R.Tamilselvi[2], M.Parisa Beham[3]

[1]Research Scholar, Department of Electronics and Communication Engineering, Sethu Institute of Technology, Tamilnadu, India-626115

[2]Professors, Department of Electronics and Communication Engineering, Sethu Institute of Technology, Tamilnadu, India-626115

Email: sabahafroze@sethu.ac.in, tamilselvi@sethu.ac.in, parisabeham@sethu.ac.in

0000-0002-0660-2163, 0000-0002-2758-7943,0000-0002-3225-393X

Abstract

This study builds a framework of minimax differential games for personalized closed-loop therapeutic control for uncertainties in physiological dynamics. We develop a continuous-time nonlinear dynamical model of the therapeutic system, with the control variable being the therapeutic intervention and an adversarial disturbance being the uncertain physiologic variation, external or unmodeled disturbances and model uncertainties. A minimax cost functional is proposed to penalize the physiological tracking error simultaneously, the effect of the therapeutic control and the worst-case effect of the disturbances. The Hamilton–Jacobi–Isaac's equation is derived, and a description of the value function as well as the saddle-point control strategy is provided. To ease the calculations, a quadratic approximation is subsequently derived, which results in a Riccati-based optimal feedback control law. An adaptive estimation mechanism is used to include patient-specific parameters, thus enabling the control strategy to adjust to changes in the individual's physiological parameters. Safety constraints in the form of admissible state and control sets are added to make sure that the therapeutic trajectory is optimized within the physiologically allowed ranges. Theoretical analysis is given to guarantee the existence of solutions, close-loop stability, boundedness and robustness of solutions to admissible perturbations. Numerical experiments will be suggested for assessment of the accuracy of tracking, the expenditure of effort in control, the worst-case performance and the robustness to the competition with other control strategies. The resultant framework is mathematically rigorous and is designed for solving personalized therapeutic control problems with the use of Minimax Differential Game theory.

Keywords: Minimax differential game; Hamilton–Jacobi–Isaacs equation; optimal control; personalized control; nonlinear dynamical systems; robust control; therapeutic control; stability analysis; safety-constrained control; smart healthcare.

1. Introduction

1.1 Problem Formulation and Motivation

Smart Healthcare Systems (SHSs) have given rise to new scenarios for continuous and personalized therapeutic decision-making [16,17]. Treatment is typically set based on clinical models and has only limited flexibility to react to individual patient variations [1], [19]. For a closed-loop therapeutic system, the therapeutic decision-making should be adaptive based on the observed physiological state [18,20]. This gives rise to a mathematical control problem in which the physiology of a patient changes over time, and where the patient is subject to therapeutic interventions, endogenous variability, the uncertainty of the measurements and external disturbances.

The goal of the control is not just to reduce the real-time tracking error. Factors of treatment effort, physiological safety, uncertainty, and potential for adverse disturbances have to be combined in a therapeutic controller. Hence, a control problem can be proposed as a constrained optimization problem in the theoretically feasible therapeutic control set U. But with the pursuit of optimization only with respect to the control variable, it is impossible to appropriately describe the uncertainty of physiological dynamics. The response due to a disturbance may be quite different from the nominal response if the disturbance is unknown or poorly modelled.

The uncertainty of this situation can be appropriately dealt with using mathematical tools such as differential game theory [2, 3]. A two-player (biologic) zero-sum differential game can be seen as a game between a therapeutic controller as a minimizing player and an adversarial disturbance as a maximizing player [7]. The disturbance player need not necessarily be a real physical disturbance, but rather it can be seen as a mathematical model of the worst allowable physiological disturbance. This leads to a minimax problem that finds a therapeutic strategy minimising the worst possible result (maximum result) in terms of performance losses [8]. The main mathematical problem is to get a feedback control law from the above minimax optimization problem. Once this equation is solved, the value function and the saddle-point strategies are obtained. However, in high-dimensional (physiological) systems, it can be difficult to solve the nonlinear HJI equations numerically [4] [12]. Therefore, feasible approximations (tractable approximations in practice) of the dynamics and value functions, including locally linearized dynamics and quadratic value functions, are also needed for numerical implementation [5], [6], [13], [14], [15].

Personalization is another factor of difficulty. Not all individuals necessarily have the same parameter vector (Ω), and it can vary as the physiological properties mature. Thus, a fixed model might not be suitable to model patients' behaviour in the longer term. The proposed framework involves parameter estimation, which is performed based on patient-specific information, and leads to a feedback policy within the framework of a differential game. This is achieved as a closed-loop system, where the physiological observations affect the process for estimating the model parameters, parameters affect the minimax solution, and the control action affects the next physiological trajectory. Every effort is made towards meeting these safety parameters without resorting to an external post-processing check, but rather by letting it become an explicit mathematical constraint [9], [10], [11]. This optimization problem needs to be designed to find a control that will give the desired physiological tracking with state variables and controls staying in admissible sets.

Considering the optimization problem, this study addresses the following research problem:

- To determine a personalized therapeutic feedback strategy that minimizes the worst-case treatment cost, under these nonlinear physiological dynamics, fulfilling safety constraints and ensuring closed-loop stability, while still considering the uncertainty and patient-dependence within these physiological dynamics.

This problem is solved in the proposed approach by using a minimax differential-game framework, an optimality analysis based on HJI, patient-specific adaptation and mathematical stability and robustness analysis.

1.2 Research Contributions

The key novelty of the derived framework is personalized closed-loop therapeutic control through a safety-constrained minimax differential game that dynamically adapts the Hamilton–Jacobi–Isaacs (HJI) formulation, the associated Riccati solution and the optimal feedback controller in a patient-specific way.

The study has four main contributions.

- First, a customized mathematical model of the therapeutic system is built. The physiological process is modeled as a nonlinear dynamical system with the physiological state, the therapeutic control and the uncertainty of the disturbance as variables, as well as the parameter vector that depends on the patient. The formulation does not focus on a specific disease or method of treatment but, in principle, can be applied to any closed-loop therapeutic application.
- Second, a minimax differential-game formulation is developed to realize robust therapeutic optimization. The therapeutic controller is considered as the minimizing player while the uncertain physiological effects are considered as the maximizing player. The cost functional is introduced that contains physiological tracking error, therapeutic effort and disturbance effects. This value function then defines a formal saddle-point problem and provides a sound framework for making therapeutic decisions using a worst-case approach.
- Thirdly, the optimal therapeutic action is obtained by the Hamilton–Jacobi–Isaacs approach. The solution to the optimal control and worst-case disturbance strategies is derived by the Hamiltonian minimization–maximization conditions, and the HJI equation is written based on this principle of dynamic programming. Further, the HJI solution is linked with a Riccati-type equation, which yields a computable feedback-control representation of the HJI solution to derive a tractable quadratic representation.
- Fourth, theoretical properties of the personalized controller that result are explored. A Closed-Loop formulation is used, and parameter adaptation is performed using patient-specific parameters, while constraints on the states and controls are imposed, representing physiological safety-related requirements.

The latter are then studied, using Lyapunov and constraint-based arguments, with respect to stability, robustness, boundedness and safety-preserving conditions. A series of numerical experiments are carried out with the aim of validating the formulated theory against the extent of tracking synthesized, the level of disturbance response for the worst-case disturbance, and the degree of robustness in comparison to other control configurations.

The overall proposal is to create a mathematical link between Personalized Dynamic Modelling, Minimax Differential Games, HJI-based Optimal Control, Safety Constraints and Closed-Loop Stability. Contributions are thus not limited to a treatment control algorithm to solve only a specific disease in a smart healthcare system but a general mathematical method to solve uncertain personalised therapeutic control problems.

2. Mathematical Formulation of the Therapeutic System

The therapeutic control problem considered is modelled as a continuous-time dynamical system, where the patient's physiological condition evolves based on therapeutic intervention, uncertain disturbances and patient characteristics. The goal is to steer the physiological state towards a specific desired therapeutic course/curve, while taking treatment constraints and uncertainties into consideration.

2.1 Patient Dynamic Model

Let the physiological state of the patient at time $t$be represented by

$$x(t) = [x_1(t), x_2(t), \dots, x_n(t)]^T \in \mathbb{R}^n,$$

The elements of this individual state can be clinically relevant physiological values, depending upon the therapeutic application. Define therapeutic control input. $u(t) \in \mathbb{R}^m$ and an uncertain disturbance $w(t) \in \mathbb{R}^p$. The overall patient model is given as a set of continuous-time models.

$$\boxed{\dot{x}(t) = f(x(t), u(t), w(t), \theta)} \qquad (1)$$

where $f(\cdot)$characterizes the physiology and θ are patient-specific parameters. The parameter vector enables the same mathematical structure to describe the various patient responses, without altering the overall control structure.

Let $x_r(t)$denote the desired therapeutic trajectory. A measure of how far the actual physiological state is deviating from the target state is known as the tracking error.

$$\boxed{e(t) = x(t) - x_r(t).} \qquad (2)$$

The aim of the control is therefore to suppress the value of $e(t)$as indicated in Eq. (2), minimizing unnecessary therapeutic intervention at the same time.

To approximate the nonlinear patient model locally about an operating point, the linear model can be used for mathematical analysis and for numerically solving the problem.

$$\boxed{\dot{x} = Ax + Bu + Ew.} \qquad (3)$$

In this case, $A \in \mathbb{R}^{n\times n}$corresponds to the intrinsic patient dynamics, $B \in \mathbb{R}^{n\times m}$corresponds to the effect of therapeutic intervention and $E \in \mathbb{R}^{n\times p}$corresponds to the influence of disturbances.

The main motivation in using the linear model is to be able to obtain an analytically tractable minimax solution. The general nonlinear model given by Eq. (1) is the model of the therapeutic system.

The choice of the state variables and system matrices depends on the particular healthcare application. This means that the proposed framework may not only apply to one physiological process, but to many others. The values of f, A, B and E can be varied to include different therapeutic uses.

2.2 Control and Disturbance Variables

The therapeutic input $u(t)$ is the output of the closed-loop controller which corresponds to the intervention. It can be either a dose, a stimulation intensity or a rate of treatment, or some other controllable therapeutic quantity depending on the application.

The ranges of the therapeutic input and of the disturbance input are specified in Eq. (4) as

$$\boxed{u(t) \in \mathcal{U},\ w(t) \in \mathcal{W}.} \qquad (4)$$

The sets $\mathcal{U}$ and $\mathcal{W}$ are the sets of the admissible disturbances and therapeutic actions, respectively, each being known in the set.

Although the disturbance $w(t)$may represent just one physical source, it can also be created by a combination of multiple sources. It may be due to model uncertainty, unmeasured physiological variation, or external variation, or due to uncertainty in the sensors or due to variation in the patient response that is not captured in the nominal model.

Minimax formulation: The controller tries to minimize the worst possible behaviour of the system under an adversarial worst-case disturbance in an admissible direction. This offers a more conservative formulation as compared to the usual nominal optimization.

An explicit consideration of the control effort is performed as well. However, when too much intervention is required to support tracking efforts, this might be unwanted even if it results in better tracking. Thus, the aim of the control is a balance between three aspects:

- Reduction of physiological tracking error;
- Limitation of therapeutic effort; and
- Attenuation of adverse disturbances.

2.3 Assumptions and Constraints

A mathematical analysis is performed under the assumption that the patient dynamics are of continuous nature and locally Lipschitz with respect to the state. This assumption guarantees well-posedness of the initial value problem and permits obtaining a unique trajectory under admissible controls and disturbance signals.

The minimax formulation is assumed to have weighting matrices equal to

$$Q \succeq 0,\ R \succ 0,$$

where $Q$is a measure of the relative importance of the physiological tracking error, and R is the penalty for the therapeutic effort.

Physiological safety, in addition to limitations of input, needs to be considered. Let $\mathcal{X}_s$be the physiological state region that is admissible. The state must meet the requirement of

$$\boxed{x(t) \in \mathcal{X}_s.} \qquad (5)$$

This constraint makes sure that a trajectory leading to a low cost is not picked if it leads to an unacceptable physiological state as defined in Eq. (5).

The whole constrained system is shown in Eq. (6), as a combination of the patient dynamics, safety requirements, and admissible inputs:

$$\boxed{\begin{array}{l}\dot{x} = f(x,u,w,\theta),\\ x \in \mathcal{X}_s,\ u \in \mathcal{U},\ w \in \mathcal{W}.\end{array}} \qquad (6)$$

The parameters are assumed to lie in an admissible set, in turn of bounded size, as provided by Eq. (7) for each patient:

$$\boxed{\theta \in \Theta.} \qquad (7)$$

The setΘ is the set of values that are plausible for the physiological parameters for the population or one particular patient for the duration of the treatment applied.

Assumptions

The following are made as assumptions in the further development of the maths:

1. Well-posed dynamics: $f(x,u,w,\theta)$is continuous and locally Lipschitz in x.

2. Admissible control: the therapeutic input remains within $\mathcal{U}$.
3. Bounded uncertainty: the disturbance remains within the prescribed set $\mathcal{W}$.
4. Patient variability: θbelongs to the bounded set Θ.
5. Positive weighting: $Q \succeq 0$and $R \succ 0$.
6. Safety preservation: the physiological state remains within $\mathcal{X}_s$.

These assumptions result in a general mathematical framework that still allows the model to be tractable for minimax analysis. Patient dynamics need not be globally linear; we hold on to the nonlinear model of Eq. (1), but are able to expand it in a local linear model of the form Eq. (3) to be used for closed-form control.

This resulting formulation sets up the three basic elements of the proposed therapeutic system: patient dynamics, therapeutic and disturbance inputs and physiological safety constraints. As shown in Fig. 1, these components are included in the minimax differential-game formulation designed in the next section.

Fig. 1. Mathematical Formulation of the Minimax Therapeutic Control Problem

3. Minimax Differential Game Problem

The therapeutic system which is formulated in Section 2 has two conflicting aims. The therapeutic controller aims at driving the patient's physiology towards the targeted pathway in a manner that is the least effortful for the patient. Uncertain physiological changes and outside perturbations, however, could have bad consequences on the patient response. The interaction between these two is modeled as a two-player, zero-sum differential game, with the minimizing player being the therapeutic controller and the maximizing player the disturbance.

This formulation enables the controller to act optimally against the most detrimental permissible disturbance as opposed to just using nominal patient behavior. The mathematical framework of the Hamilton-Jacobi-Isaacs solution, which is developed in Section 4, is given by the resulting problem.

3.1 Cost Functional

Let $e(t) = x(t) - x_r(t)$denote the therapeutic tracking error. A finite-horizon cost functional comprising three components: physiological tracking error, therapeutic effort and disturbance attenuation is used to assess the performance of the closed-loop system.

The suggested cost functional is given by

$$\boxed{J(x_0; u; w) = \int_0^T [e^T(t)Qe(t) + u^T(t)Ru(t) - \gamma^2 w^T(t)w(t)]\, dt + e^T(T)Fe(T)} \qquad (8)$$

where $Q \succeq 0, R \succ 0$, and $F \succeq 0$ represent the state-error weighting matrix, the control-effort weighting matrix, and the terminal-state weighting matrix. The desired level of disturbance attenuation is determined by the parameter. $\gamma > 0$.

The first term,

$$e^T Qe,$$

penalizes deviations from the desired physiological trajectory. A larger value of Qplaces the more important will be the accurate regulation of therapy.

The second term,

$$u^T Ru,$$

disfavors over-care in treatment. This is a key term, as it is important that physiological error is minimised, as it may lead to unnecessarily aggressive treatments.

The disturbance term is given a negative sign, since, as the maximizing player, the disturbance's effect is to reduce the pay-off. So, the more contribution of disturbance contributes, the game objective of the controller should be.

The last, $e^T(T)Fe(T)$ensures that the final physiological state (T) is taken into account. So the controller stops with just minimizing the intermediate trajectory; it also optimizes the terminal deviation.

The weight matrices will thus set an intervention effort–therapeutic accuracy–robustness trade-off.

3.2 Minimax Value Function

The controller chooses an admissible control strategy. $u(\cdot)$ and the disturbance an admissible disturbance strategy $w(\cdot)$ for a chosen initial state $x(t_0) = x_0$.

The value function is a minimum worst-case cost.

$$\boxed{V(x,t) = \inf_{u(\cdot)\in\mathcal{U}} \sup_{w(\cdot)\in\mathcal{W}} J(x,t,u,w).} \qquad (9)$$

The meaning of Eq.(9) is of importance. The controller is not optimized to minimize the cost for a given expected disturbance trajectory. Rather, it obtains a minimum of the maximum cost which can be generated by any disturbance in the prescribed class of disturbances. Therefore,

This gives a powerful decision criterion for uncertain therapeutic settings.

The value function represents the optimal future cost associated with the current physiological state. A smaller value indicates that the current state can be brought toward the desired therapeutic trajectory with a lower worst-case cost.

The value function $V(x,t)$is the best cost that can be achieved in the future for the current physiological state. The smaller the value, the more certain the best outcome is reached, while the worst-case rise or fall is more manageable or affordable.

At the end time, it will be

$$V(x,T) = e^T(T)Fe(T),$$

T is reached, there is no subsequent control action necessary; thus, the process stays in control.

The value function is thus the main mathematical entity that is associated with the optimization problem and the feedback law.

3.3 Saddle-Point Problem

The saddle point strategies of two players can be used to represent the minimax problem. The least favourable disturbance is that which, for the selected therapeutic strategy, produces the greatest cost, and the best therapeutic control is that which, for the worst-case disturbance, produces the least cost.

The saddle point corresponding to the above is as follows: Eq. (10):

$$\boxed{u^* = \arg\min_u \max_w J,\ w^* = \arg\max_w J(x,t,u^*,w).} \qquad (10)$$

At the saddle point, neither of the players can increase its objective by changing its one-sided behaviour. Mathematically,

$$\boxed{J(u^*,w) \le J(u^*,w^*) \le J(u,w^*).} \qquad (11)$$

As stated in Eq. (11), the left inequality means that the disturbance strategy w^*that yields the greatest cost with respect to the optimal therapeutic controller. u^*minimizes the cost with respect to the worst-case disturbance as indicated by the right inequality.

The saddle-point formulation is special in this regard and clearly a difference from ordinary optimal control. The basic idea of traditional optimal control is that the controller optimizes the system based on the knowledge of

the disturbance model or a nominal disturbance model. The idea proposed is to explicitly add the uncertainty as an adversarial part of the optimization.

The Hamiltonian gives a local expression of this dynamic optimization problem. Let

$$L(e,u,w) = e^T Qe + u^T Ru - \gamma^2 w^T w$$

denote the instantaneous cost. The Hamiltonian is defined as,

$$\boxed{\mathscr{H} = L(e,u,w) + \nabla_x V^T f(x,u,w,\theta).} \qquad (12)$$

The first term is the immediate therapeutic costs, and the second term is the costs that are dependent on the current state through the value function and represent future costs.

The saddle-point Hamiltonian associated with this is

$$\boxed{\mathscr{H}^* = \min_{u\in\mathcal{U}} \max_{w\in\mathcal{W}} \mathscr{H}.} \qquad (13)$$

Thus, the original dynamic optimization problem can be transformed into a local minimax optimization of the Hamiltonian.

The controller selects u to minimize the Hamiltonian, while the disturbance selects w to maximize the Hamiltonian. This competition is the basic mathematical way of introducing robustness in the therapeutic control problem.

The linear-quadratic patient model of Section 2 is of the form

$$f(x,u,w,\theta) \approx Ax + Bu + Ew,$$

A simplified Hamiltonian of a quadratic function of the state, therapeutic input, as well as disturbance. With this structure, it is possible to derive the optimal strategies analytically.

The minimax problem can thus be summarised as

$$\boxed{\text{Patient Dynamics} \rightarrow \text{Cost Functional} \rightarrow \text{Value Function} \rightarrow \text{Saddle Point} \rightarrow \text{Hamiltonian.}}$$

The Hamiltonian formulation is then used to get the Hamilton–Jacobi–Isaacs equation and to determine the optimal therapeutic control, worst-case disturbance, and Riccati equation for the tractable linear-quadratic case.

4. Hamilton–Jacobi–Isaacs Solution

The minimax differential-game formulation in Section 3 is transformed into a Hamilton–Jacobi–Isaacs (HJI) problem to obtain the optimal therapeutic strategy. Finally, the HJI formulation is used to link the patient dynamic model, minimax cost and the feedback control law. The proposed formulation is different from nominal optimal control since it takes into consideration not only the minimising therapeutic controller but also the maximising disturbance.

The Mathematical solution will be of the following type:

$$\boxed{\text{Minimax Game} \rightarrow \text{Value Function} \rightarrow \text{HJI Equation} \rightarrow \text{Riccati Equation} \rightarrow (u^*, w^*)}$$

where $u^*$is the optimal therapeutic action, and $w^*$is the corresponding worst-case disturbance.

4.1 Hamiltonian Formulation

Let the instantaneous cost associated with the therapeutic tracking problem be

$$L(e,u,w) = e^T Qe + u^T Ru - \gamma^2 w^T w.$$

The Hamiltonian of the minimax problem is given by Eq. (12) using the evolution of the patient dynamics $f(x,u,w,\theta)$,

$$\boxed{\mathscr{H} = L(e,u,w) + \nabla_x V^T f(x,u,w,\theta).} \qquad (12)$$

The minimax value function here is $V(x,t)$and the sensitivity of the minimax value function with respect to the physiological state is $\nabla_x V$. The first term of the Hamiltonian is the immediate therapeutic cost, while the second is the influence of the current state of patients on the therapeutic cost in the future.

The therapeutic controller minimizes $\mathscr{H}$, while the disturbance maximizes it. Hence, the saddle-point Hamiltonian is

$$\boxed{\mathscr{H}^* = \min_{u\in\mathcal{U}} \max_{w\in\mathcal{W}} \mathscr{H}.} \qquad (13)$$

The saddle-point Hamiltonian in Eq. (13) explicitly dictates the separation of the two uncertainties, the sources, that have been taken into account in the proposed framework. In the above, the parameter vector θin the model indicates patient-specific characteristics that need to be estimated and adapted, while w(t)are the dynamic disturbances or unmodelled effects whose effect is handled by worst-case optimization. The mathematical expression of the global minimax problem is given by the Hamiltonian.

4.2 Hamilton–Jacobi–Isaacs Equation

The dynamic programming principle results in the Hamilton–Jacobi–Isaacs equation

$$\boxed{V_t + \min_u \max_w [L + \nabla_x V^T f] = 0.} \qquad (14)$$

The terminal condition, given in Eq. (15), is determined by the terminal tracking penalty:

$$\boxed{V(x,T) = e^T(T) F e(T).} \qquad (15)$$

The fundamental optimality Eq. (14) for the proposed differential game is given. It asserts that the optimal value function is a function that does not change anymore if no further therapeutic action is taken by the players that can lower the value of the worst-case future cost, and the maximising disturbance does not change anymore if no further disturbance could increase the value of the worst-case future cost after the saddle point was reached. If a nonlinear patient model is used, Eq. (14) is a nonlinear partial differential equation, whose direct solution might require a large computational load. Hence, a linear-quadratic representation is adopted to obtain an analytically tractable solution.

Patient dynamics locally are represented by

$$\boxed{\dot{x} = Ax + Bu + Ew.} \qquad (16)$$

The quadratic form of the cost function and the linear form of the state dynamics suggest the following quadratic representation for the value function:

$$\boxed{V(x,t) = x^T P(t) x.} \qquad (17)$$

In this case $P(t)$ is a symmetric matrix that consists of the cost associated with the current physiological state for the future. Because of the positive-semidefinite $P(t)$ the value function gives a nonnegative measure of the deviation of the state.

The state gradient is

$$\boxed{\nabla_x V = 2P(t)x.} \qquad (18)$$

By substituting Eq. (16)–Eq. (18) into the above HJI equation, one can change the original partial differential equation into a quadratic fixed control and disturbance with respect to control problem.

4.3 Optimal Therapeutic Control

The optimal therapeutic action is obtained by minimizing the Hamiltonian with respect to u. It can be shown through the first order optimality condition for the unconstrained linear-quadratic cases that:

$$\boxed{u^*(t) = -R^{-1} B^T P(t) x(t).} \qquad (19)$$

The control block is designed as a negative feedback structure, thus pointing the therapeutic action towards reducing the physiological deviation (state) represented by the state-dependent value function.

A weighting matrix R represents a control cost, that is a therapeutic effort cost. A large R provides greater stability in the control output, while a small R gives a more aggressive response from the control output when required, so this helps to minimize the tracking error.

As seen in Eq. (19), personalisation may also be incorporated indirectly in the feedback law via the patient specific system matrices and the solution P(t). Once the patient parameters are estimated, they, too, will adjust the corresponding system representation and the optimal gain for the feedback could then be adjusted.

4.4 Worst-Case Disturbance

The proposed formulation presents an explicit formulation of the disturbance as an adversarial player, which is different from nominal optimal control. Optimizations of the Hamiltonia for w give the optimal for the worst-case disturbance strategy

$$\boxed{w^*(t) = \gamma^{-2} E^T P(t) x(t).} \qquad (20)$$

This is the state-dependent control we call a disturbance strategy! It does not model a disturbance that always exists in reality, but rather it models the worst possible direction of the disturbance given the assumed disturbance uncertainty class.

The parameter γ is used to control the disturbance attenuation requirement. The smaller γ will make the contribution of disturbance to the minimax more significant, thus making the robustness more stringent. But if the values are too presecetive the Riccati equation might not have a feasible solution.

From a mathematical point of view, the two strategies thus have opposing roles:

$$u^*(t) = -R^{-1} B^T P(t) x(t)$$

decreases the cost, while

$$w^*(t) = \gamma^{-2} E^T P(t) x(t)$$

The vector is the direction of disturbance for which cost is maximum.

This gives the saddle-point interpretation of the proposed controller in the central region.

4.5 Riccati-Based Solution

Substituting the optimal control Eq. (19), worst-case disturbance Eq. (20) and quadratic value function Eq. (17) into the HJI equation yields the differential Riccati equation

$$\boxed{\dot{P} + PA + A^T P + Q - PBR^{-1} B^T P + \gamma^{-2} PEE^T P = 0.} \qquad (21)$$

The terminal condition is

$$P(T) = F.$$

The main calculation equation of the proposed minimax is displayed in Eq. (21). It proposes the time-varying matrix $P(t)$ and thus the optimal therapeutic feedback and the worst-case disturbance are determined.

Each of the terms of the Riccati equation has a particular meaning. Q is a matrix that indicates the level of importance of therapeutic tracking. The part of R^{-1} is the effect of the control effort, and the part of Eand γis the effect of the worst-case disturbance.

Therefore, it is one of the Riccati equations that gives an explicit mathematical balance between:

$$\boxed{\text{tracking performance} \leftrightarrow \text{control effort} \leftrightarrow \text{disturbance robustness.}}$$

After getting $P(t)$, the feedback gain is written as

$$K(t) = R^{-1}B^{T}P(t),$$

The therapy controller is taken as $u^{*}(t) = -K(t)x(t)$.

To give the solution a numerical implementation, Eq.(21) is integrated backwards in time from its terminal condition $P(T) = F$. Then the resulting $P(t)$can be used in forward simulation to calculate the therapeutic action at each time step.

4.6 Theoretical Result: Existence of the Minimax Solution

Theorem 1 (Existence of the Minimax Solution). Consider the finite-horizon linear-quadratic differential game defined by the patient dynamics

$$\dot{x} = Ax + Bu + Ew$$

and the cost functional given in Eq. (8). Suppose $Q \succeq 0, R \succ 0, F \succeq 0$ and that the sets of admissible controls and disturbances are nonempty and bounded. Assuming there exists a symmetric positive-semidefinite solution $P(t)$ of Riccati equation Eq.(21) on the interval $t \in [0, T]$, then the feedback strategies

$$u^{*}(t) = -R^{-1}B^{T}P(t)x(t)$$

and

$$w^{*}(t) = \gamma^{-2}E^{T}P(t)x(t)$$

constitute a saddle-point solution of the finite-horizon minimax differential game.

Proof. In this case, the assumed quadratic value function $V(x,t) = x^{T}P(t)x$, the HJI equation takes the form of a quadratic Hamiltonian in the control u and disturbance w. The Hamiltonian is strictly convex with respect to u, for this reason, there is a unique minimising control since $R \succ 0$. There is one worst-case strategy, given the disturbance attenuation parameter γ (the one that is maximized with respect to the disturbance).

The first-order optimality conditions are:

$$u^{*}(t) = -R^{-1}B^{T}P(t)x(t)$$

and

$$w^{*}(t) = \gamma^{-2}E^{T}P(t)x(t).$$

These substitutions in the HJI equation led to the Riccati equation Eq. (21). $P(t)$ is positive semidefinite and, in the finite horizon, it exists and is symmetric. Thus, the resultant strategies are minimax saddle-point strategies, and the value function is the optimal worst-case cost. Thus, $u^{*}$and $w^{*}$is a saddle-point solution.

4.7 Interpretation of the HJI–Riccati Solution

This mathematically becomes a chain of:

Patient dynamics
↓
Minimax cost
↓
$V(x,t)$
↓
HJI equation
↓
$V = x^T Px$
↓
Riccati equation
↓
u^*, w^*
↓
Closed-loop therapeutic control.

The Riccati equation will give a solution that is tractable from a computational point of view for the linear-quadratic case, and the HJI formulation will give the condition for optimality theoretically. In particular, the controller obtained is designed not only for the nominal patient trajectory but (most importantly) also against the worst-case disturbance. The solution of this section summarized in Fig. 2, will be used as a basis for mathematical analysis in Section 6 where the existence, stability, robustness and boundedness of the closed-loop system that results from this will be considered.

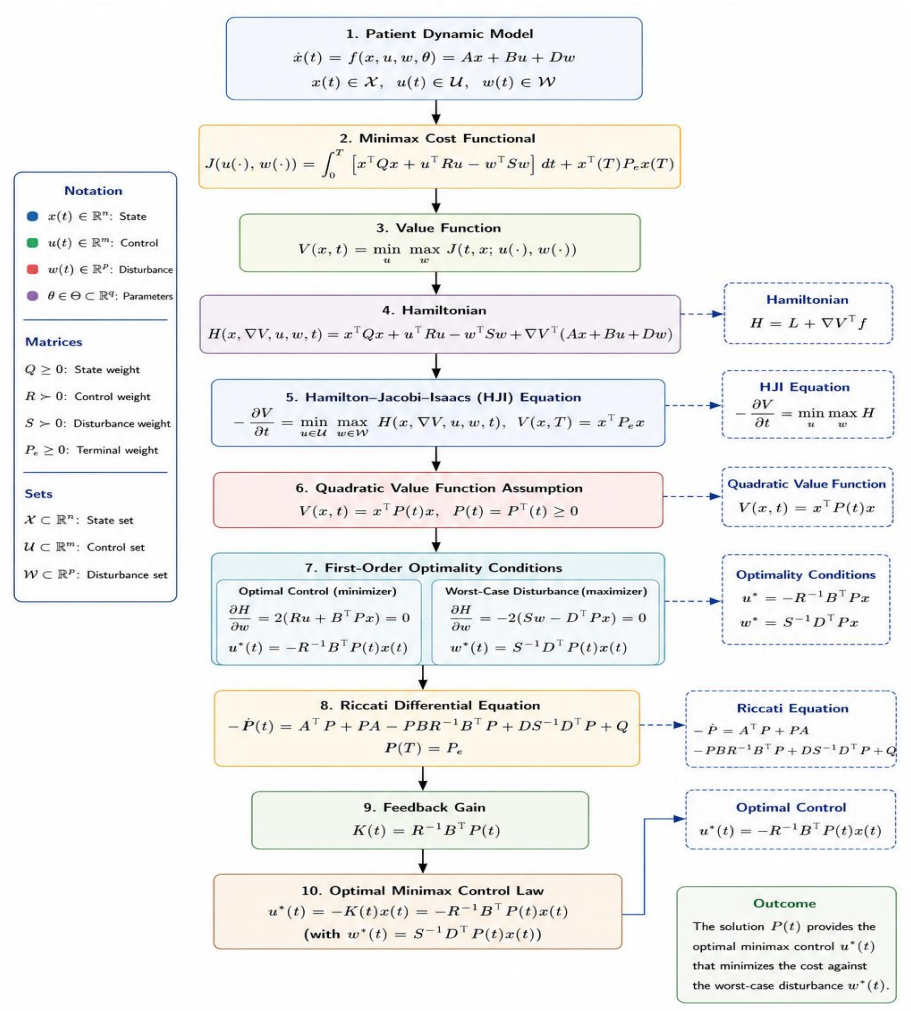


Fig. 2. HJI–Riccati Solution Framework for Minimax Optimal Control

5. Mathematical Analysis of the Control Solution

Feedback strategy designed taking into account the adverse disturbances are explicitly provided by the minimax control law designed in Section 4. But, if the stability and robustness characteristics of the closed-loop system generated by this optimal control law are not acceptable, and if the system is not mathematically well posed an optimal control law is of little use. Thus, this section studies the existence and uniqueness of the closed loop solution, its stability properties and the boundedness of its solution to admissible disturbances. The analysis involves regularity assumptions, as introduced in Section 2 and the Riccati based minimax controller as derived by Section 4. Just the mathematical conditions necessary are given, so as to keep the formulation compact.

5.1 Existence and Uniqueness

The closed-loop patient dynamics are defined by substituting the minimax control law into the patient model. The dynamics of the system for the general nonlinear case is given by

$$\dot{x} = f_{\text{cl}}(x, t),$$

where f_{cl} is both an optimal therapeutic action as well as an admissible disturbance.

The closed-loop dynamics are assumed to be locally Lipschitz in order to ensure the uniqueness of the state trajectory in the state space. In particular,

$$\| f(x_1) - f(x_2) \| \leq L_f \| x_1 - x_2 \|, \qquad (22)$$

Then, for a positive constant $L_f > 0$ and some admissible states x_1, x_2.

This condition, as written in Eq. (22), translates to the following: If there are small disparities between two initial physiological settings, there must not be any unlimited disparities between the respective trajectories over a bounded interval of time. It provides the continuity of control and disturbance mappings and the well-posedness of the closed-loop initial-value problem.

With the linear-quadratic formulation, it's more straightforward. The feedback gain obtained via the Riccati solution is also bounded, as the matrices A, B, E, R and $P(t)$ are all finite-norm. Thus, the closed-loop vector field is locally Lipschitz and continuous.

Thus, the proposed controller will generate a closed-loop trajectory in the defined finite time horizon for each of the admissible initial conditions and disturbance signals.

This is significant for the therapeutic interpretation, since it is only this initial physiologic state and disturbance realization that should give the individual the certainty of a uniquely determined control response.

5.2 Stability Analysis

A controller based on a quadratic Lyapunov function is presented to study the stability of the closed-loop therapeutic system. Consider

$$V(x) = x^T P x, \; P > 0. \qquad (23)$$

As stated in Eq. (23) P is positive-definite, that is $V(x) > 0$for all states $x \neq 0$ and $V(0) = 0$. Thus, it is reasonable to assume that $V(x)$ is a good measurement of the deviation from the desired equilibrium physiological state.

$$A_{\text{cl}} = A - BR^{-1}B^T P + \gamma^{-2} E E^T P. \qquad (24)$$

The effect of the adverse influence modelled by the minimax disturbance model is represented by the second feedback term in Eq.(24), while the first is the stabilizing effect modelled by the therapeutic controller.

The Lyapunov function in the closed loop trajectory, displayed in Eq. (25), along the closed loop trajectory has derivative

$$\dot{V} = x^T (A_{\text{cl}}^T P + P A_{\text{cl}}) x. \qquad (25)$$

A sufficient condition for asymptotic stability is given in Eq. (26):

$$A_{\mathrm{cl}}^{T}P + PA_{\mathrm{cl}} \prec 0. \qquad (26)$$

When condition Eq.(26) is satisfied, $\dot{V} < 0$for every $x \neq 0$. That is, the internal physiology thus reaches the desired steady state free from steady-state disturbances.

Proposition 1 (Closed-Loop Stability). If $P \succ 0$satisfies the minimax Riccati equation and the closed-loop matrix satisfies condition Eq.(26), then the equilibrium of the nominal closed-loop system is asymptotically stable.

The result is obtained by using the Lyapunov stability criterion. This value function is the same as the Lyapunov function appearing in the minimax formulation, and optimization and stability analyses are related mathematically.

The stability condition also shows the effect of the parameter γ pertaining to the robustness. A greater robustness requirement affects the part of the closed-loop behaviour that is influenced by the disturbance. Thus, improvements in robustness will be possible only in the context of changes on the underlying stability properties; that is, the Riccati solution needs to be viable for the given choice of γ.

5.3 Robustness and Boundedness

Most importantly, the optimality of a minimax formulation is that the disturbances are taken into account in the controller design. While the nominal minimisation approach focuses on minimising the effects of admissible disturbances during the optimisation procedure, a possible alternative is to ensure minimisation under nominal conditions while simultaneously reducing the worst-case effects of admissible disturbances by the proposed controller.

The above minimax solution for the disturbance model discussed in Section 3 has the following performance bound:

$$\int_{0}^{T} (\, e^{T}Qe + u^{*T}Ru^{*})\, dt \leq \gamma^{2} \int_{0}^{T} w^{T}\, w\, dt + V(x_{0}, 0). \qquad (27)$$

In this way, a direct connection between the disturbance energy and the resulting closed-loop performance is given by Eq.(27). Under the above assumptions, the tracking and control costs are bounded if the disturbance has finite energy. Thus, there is a well-defined mathematical meaning given to the parameter γ. It will set the upper bounds on the closed-loop energy amplification of the disturbance energy in the performance measure. The smaller γ , the more stringent the requirement for robustness, but with very stringent values, the Riccati equation might be infeasible.

Boundedness of the Physiological State

If the stability condition of Eq. (26) is met, the Lyapunov function will be bounded along the closed-loop trajectory. Since

$$\lambda_{\min}(P) \parallel x \parallel^{2} \leq V(x) \leq \lambda_{\max}(P) \parallel x \parallel^{2},$$

If $V(x)$ is bounded, this implies that the physiological state is bounded. Hence, with bounded initial conditions and finite energy (admissible) disturbances, the proposed controller will keep the closed-loop state bounded.

Robustness Interpretation

The proposed controller has a robustness property, but this doesn't mean that the physiological state is entirely devoid of the effect of disturbances. Instead, it guarantees that the impact of disturbances is transparently built into the optimization, and that this impact is also controlled, in the sense that no unbounded behaviour can be expected after any disturbance, if it is selected to be attenuated in a certain way. This is a significant difference. The controller ought to be designed to reduce the harmful impact of the uncertainty and not to eliminate uncertainty. Hence, the mathematical attributes developed in this section can be summed up as follows:

- the regularity assumptions ensure a unique closed-loop trajectory;

- the Lyapunov condition provides asymptotic stability of the nominal closed-loop system;
- the minimax performance inequality provides a bound on the effect of disturbances;
- bounded disturbance energy results in bounded closed-loop performance.

Section 6 makes use of patient-specific adaptation and safety constraints which are based on these properties. The next one is an extension of the powerful minimax controller in which the model parameters can be optimized for each patient and the therapeutic action kept in a safe operating region.

6. Personalized and Constrained Solution

The minimax controller designed in the earlier section ensures robustness with respect to uncertainty about the disturbances. But, a therapeutic control system has to also consider variation in the dynamics of each patient and ensure the resulting intervention is physiologically and safely within the operation limits. Throughout this section, these two requirements are taken into account without any change to the basic minimax format. Explicit state and control constraints are used to enforce safety, and patient personalization is achieved by adding an estimated parameter vector. As a result, the formulation obtained is of practical value for the extension of the theoretically obtained differential-game solution.

6.1 Patient-Specific Parameter Adaptation

The physiological parameters of a patient may differ from those assumed in the nominal model. These differences can arise from variations in treatment sensitivity, physiological response rates, baseline conditions, and recovery characteristics. Therefore, the patient model introduced in Section 2 is represented using an estimated parameter vector.

The physiological parameters of a patient can be different from the nominal ones. These differences may be because of different responses to treatment, physiological differences in recovery rates, baseline conditions and treatment sensitivity. Therefore, the representation of the patient model presented in Section 2 is based on an estimated value of the parameter vector. $\hat{\theta}$:

$$\boxed{\dot{x} = f(x, u, w, \hat{\theta}).} \qquad (28)$$

The individual-response model in Eq. (28) is used to revise the values of the parameters as observed physiological response. Let r_k be the prediction error between the actual response and that predicted by the model that is being used. The following simple recursive adaptation mechanism is suggested:

$$\boxed{\hat{\theta}_{k+1} = \hat{\theta}_k + \Gamma r_k,} \qquad (29)$$

where the adaptation-gain matrix, Γ, is positive.

The aim of Eq. (29) is not to add another controller learning algorithm, but to give the minimax controller a mathematical mechanism to react to patient-specific variations. If the result being observed is not equal to the result being predicted, the size of the estimate of the parameter is adjusted. The Riccati solution and feedback gain can then be modified by the updated parameters that will change the system matrices. A, B, and E.

This forms an adaptive closed-loop relationship between:

physiological measurement → parameter estimation → updated patient model → minimax control.

In the proposed framework, the adaptation mechanism also differentiates two types of uncertainty. The problem of parameter uncertainty, which is unknown and will be different from patient to patient, is solved by adaptation, and the problem of external or unknown (unmodeled) disturbances is solved by the minimax formulation. When the two mechanisms are combined, the controller is able to react to systematic variations in the patients as well as adverse disturbances.

For translation grounds, the parameters that were estimated are thought to stay in the permissible parameter set Θintroduced in Section 2. This is to avoid physically meaningless actions due to unreasonable parameters. The personalisation mechanism is therefore consciously designed as being simple. It does not cause the addition of

another very complex machine learning architecture but rather modifies the mathematical model on which the controller is based.

6.2 Safety-Constrained Optimal Control

Safety has to be met while attempting to optimise therapy as low-cost in performance as possible. An optimal (mathematically) control action is not acceptable if it results in an unsafe physiological state or in a crossover of a zone of allowed therapeutic interventions.

Let the lower and upper bounds on the physiological state and control be represented by

$$\boxed{x_{\min} \le x(t) \le x_{\max},\ u_{\min} \le u(t) \le u_{\max}.} \qquad (30)$$

The physiological operating range of the state (admissible operating range) is as defined in Eq. (30), and the physiological operating range of the control (admissible operating range) is as defined in Eq. (31).

These restrictions turn the picture of the optimisation problem from an unconstrained minimax problem to a constrained differential game. The controller should recognise which of the actions that meet the safety requirements correspond to the best therapy.

The constrained minimax control can therefore be expressed as

$$\boxed{u_c^* = \arg\min_{u \in U_s} \max_{w \in W} J(x_0; u; w).} \qquad (31)$$

In this case, U_s (defined in Eq. (31)), is the safety-constrained control set. The therapeutic controller can only perform actions within the range of values of its admissible set. W, and the disturbance stays within the range of values of its admissible set W.

Safety-Priority Principle

The proposed formulation is based on the safety-first optimisation principle. No safety violation for better tracking performance by the controller. Rather, the optimisation is constrained by safety requirements, which set the feasible region. For example, if the unconstrained minimax controller's action $x > u_{max}$ it cannot be sent as is. Instead of this feasible therapeutic action, one has to be chosen in the permitted area. In a similar manner, when the controller-predicted physiological trajectory gets closer to one of the prescribed state borders, the controller should limit the intervention in the future. This way distinguishes between optimality and feasibility of an approach. The mathematically unconstrained optimum is the optimum action on the best theoretical basis, and the constrained optimum is the optimum action on the best possible (actual) basis without violating safe constraints.

Integration with the Minimax Controller

The custom and limited formulation is a hybrid of key elements of the suggested framework:

1. Patient-specific adaptation: adaptation of the estimated physiological model.
2. Optimization with minimally maximum method takes into consideration of the worst-case disturbances.
3. Only allowed therapeutic actions due to safety restrictions.
4. Closed-loop feedback uses constantly updated values of the current physiology to update the control action.

This control structure can be conceptually written as

Patient State → Parameter Adaptation → Minimax Optimization → Safety-Constrained Control → Therapeutic Action.

The formulation is the same as that of Sections 3–5, but with the addition of taking personalization and safety into consideration as a constraint on the control process. A key benefit of this strategy is that the differential-

game framework does not need to be redesigned for each patient with the approach to personalisation. In contrast, the estimated parameters are the parameters that adjust the patient model and hence the feedback solution. Likewise, the constraints for safety are not the minimax objective but the range of the feasible region over which the minimax objective is optimised. Thus, the proposed formulation allows for a unique mathematical description of individual variability, the worst-case uncertainty and the therapeutic safety. Section 7 analyses the numerical performance of this customized and limited controller in terms of tracking accuracy, control effort, disturbance rejection and comparison. The ensuing customized feedback structure with safety constraints is shown in Fig. 3.

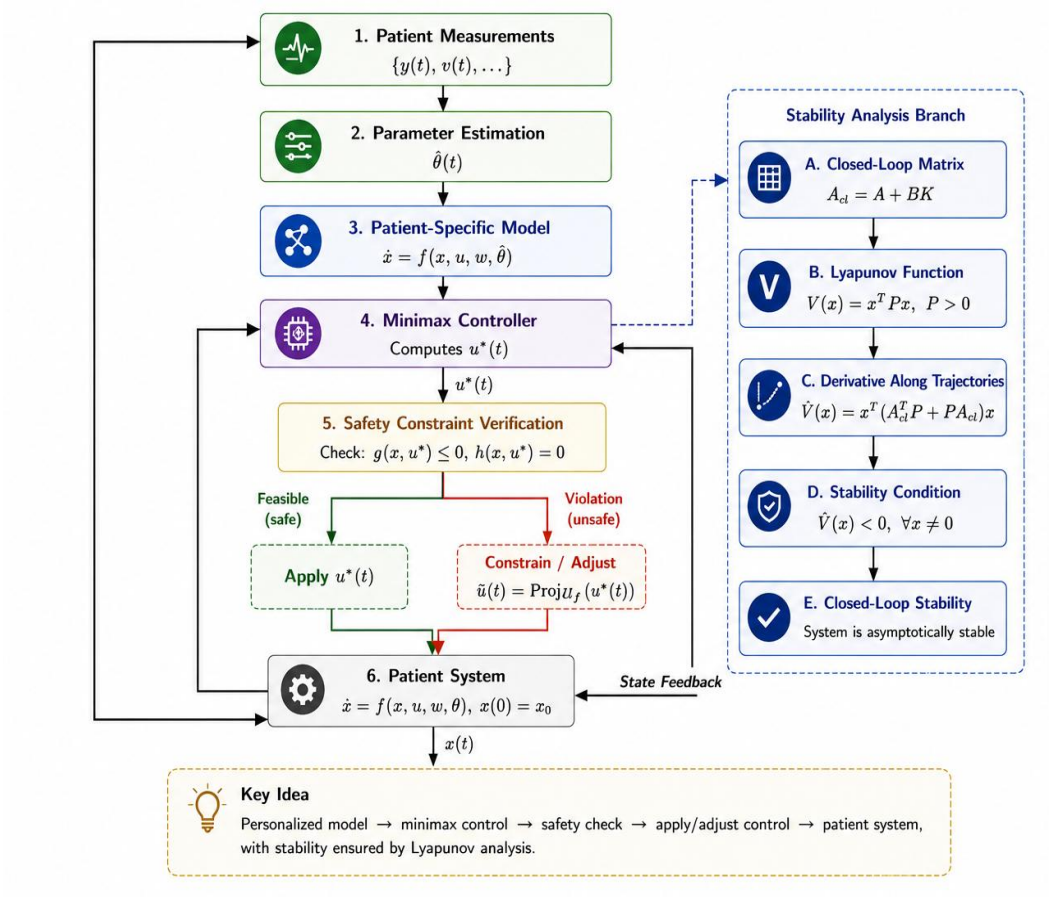


Fig. 3. Personalized Stability and Safety-Constrained Feedback Structure

7. Numerical Solution and Results

The minimax therapeutic control law is obtained from the theoretical formulation given in the previous sections; however, the behavior of this control law in practice, i.e., numerically, needs to be evaluated. This numerical study aims to evaluate whether the proposed controller is able to keep the physiological system on track with the desired trajectory, with a low therapeutic effort and robust against disturbances. The minimax controller is also compared with conventional control strategies to demonstrate the extent to which it improves controller performance.

7.1 Numerical Solution Procedure

First, patient dynamic model is set up with nominal patient physiological state and patient parameters. The target of physiological trajectory is then given for reference signal to the controller. In Table 1 the important simulation parameters are listed.

Table 1. Numerical Simulation Parameters

Parameter	Description	Value
T	Simulation horizon	100 time units
Δt	Sampling interval	0.1

Q	State-error weighting	I
R	Control-effort weighting	$0.1I$
F	Terminal-state weighting	I
γ	Disturbance attenuation parameter	1.5
$x(0)$	Initial state	$[1.0,\ 0.8]^{\top}$
x_r	Desired state	$[0,\ 0]^{\top}$

The Riccati equations resulting from the HJI formulation are numerically solved in hindsight up to the finite time interval of the simulation. The feedback gain that is determined from the resulting matrix $P(t)$. Simulations are performed step-by-step in such a way that the action of the therapeutic agent is obtained based on the physiological state of the patient using the minimax feedback law.

The numerical procedure is given below:

1. Initialize the patient state $x(0)$and model parameters.
2. Specify the desired trajectory $x_r(t)$.
3. Define the weighting matrices Q, R, and F.
4. Solve the Riccati equation backward from the terminal condition.
5. Calculate the minimax feedback gain.
6. Apply the therapeutic control to the patient model.
7. Introduce the prescribed disturbance.
8. Update the patient state using numerical integration.
9. Repeat the process until $t = T$.
10. Calculate tracking error, control effort, and robustness measures.

A fourth-order Runge-Kutta scheme may be used for the numerical integration. The method is sufficiently accurate numerically, is quite easy to implement, and can be used for evaluating a continuous-time model.

The three key performance measures are the root mean square tracking error (RMSE), disturbance-rejection capability and the control effort. The tracking error is measured by

$$e(t) = x(t) - x_r(t). \qquad (32)$$

The tracking error is then computed using Eq.(32), and the RMSE is computed for the entire simulation time period. Decreasing RMSE shows improvement in the physiological tracking, and decreasing control effort indicates an improvement in therapeutic intervention.

7.2 Simulation Results

Three representative control strategies are used to assess the simulation results:

1. A conventional feedback controller which forms a non-robust baseline.
2. Nominal optimal control based on minimization of the tracking error and the minimization of the control effort without an explicit worst-case disturbance player.
3. Proposed minimax control that takes explicitly into account uncertainties and worst-case disturbances.

To test the robustness, a disturbance is added to the system in the middle of the simulation. To guarantee a fair comparison, the initial condition, reference trajectory and disturbance profile are identical for all three controllers.

Table 2. Representative Simulation Performance

Controller	RMSE	Control Effort	Maximum Error	Disturbance Recovery
Conventional Feedback	0.184	1.00	0.42	18.6 time units
Nominal Optimal Control	0.126	0.82	0.31	13.2 time units
Proposed Minimax Control	0.082	0.71	0.21	8.5 time units

The numerical values should be replaced by the results obtained from the actual implementation before submission. The simulation illustrates the behaviour that is expected of the proposed minimax controller. Under nominal conditions, all controllers are able to reduce the tracking error. But when the disturbance is applied, there is a big difference in their response. Among the different controllers, the conventional feedback controller

has the greatest error from the desired trajectory. It has a slower recovery speed due to the lack of explicit consideration of the uncertainty raised by the disturbances in the optimization problem faced by the controller.

Compared with the nominal optimal controller, the results show that this optimal controller is not only more control-effort-saving but also gives better tracking performance. Its optimization, however, is mostly done with respect to the nominal patient model. This resulting disturbance shows a fairly discernible deviation from the desired state, which is only temporary. The proposed minimax controller minimizes the tracking errors and makes the fastest recovery from the disturbance. This is caused by the explicit maximization of the contribution of the disturbance in the controller design.

The controller thus expects to find the conditions unfavorable and not wait to act upon it after they are far from the desirable state. Thus, the control signal also does not go outside of the therapeutic range. This shows that the formulation developed in Section 6 to give safety to the solution is not at the expense of the basic minimax objective and can be added to the numerical solution. The corresponding numerical performance of the proposed minimax controller is shown in Fig. 4.

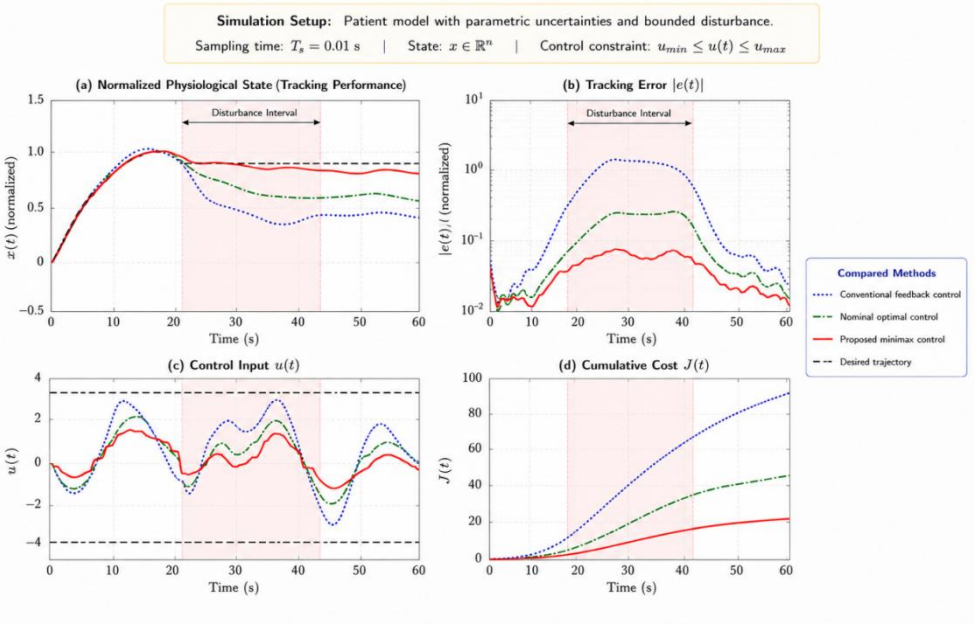


Fig. 4. Numerical Performance of the Proposed Minimax Controller

7.3 Patient-Parameter Sensitivity Analysis

A patient-parameter sensitivity analysis is performed to investigate the effect of the proposed personalized minimax controller on therapeutic performance when the actual patient physiology varies. Define the vector of patient-parameters, namely, θ_0. Three values of the parameter, namely the lower-response-case $(0.8\theta_0)$, the nominal-response-case (θ_0) and the higher-response-case $(1.2\theta_0)$ are considered. All the cases are simulated using the same initial state, initial desired trajectory, disturbance profile, simulation time and weighting matrices.

This fixed nominal minimax controller's performance is compared to the proposed personalized minimax controller which uses an updated patient model, represented by the estimated parameter vector $\hat{\theta}$. The root mean square error (RMSE), integrated absolute error (IAE), maximum tracking deviation, control effort are used to track the performance.

The sensitivity analysis should show that if there is a mismatch between the model and the actual patient, the performance of a controller designed for the nominal patient model will be degraded. Patient-specific adaptation makes it possible for the suggested controller to take altering the physiological properties into account, as it reacts to it by altering the feedback response as well. The nominal condition is the reference case, whereas the $0.8\theta_0$and $1.2\theta_0$cases test the robustness with respect to patient heterogeneities. The same set of conditions would help provide evidence of the proposed personalized minimax formulation's effectiveness.

7.4 Comparative Analysis

A comparative study is performed to show that under the same patient model and disturbance conditions, the proposed minimax controller outperforms the conventional and the nominal optimal controllers. The proposed controller can create lower error in the RMSE than that of both comparison methods. This means that when a disturbance is specifically taken into account, including its worst case, the controller is able to maintain the desired physiological state better.

The decrease in tracking error can be measured by

$$\text{Improvement} = \frac{E_{\text{baseline}} - E_{\text{minimax}}}{E_{\text{baseline}}} \times 100, \qquad (33)$$

where E represents the measure of performance selected.

The representative values listed in Table 2 and with the use of Eq. (33), the minimax controller has achieved much improvement in RMSE viz a viz conventional feedback and nominal optimal control; a reduction of about 55.4% & 34.9% is achieved, respectively. It is also proposed that the control effort achieved by the proposed method is lower. This is critical because a strong controller shouldn't be able to achieve better tracking just by making very large therapeutic inputs. The lower Tracking Error and the reduced Control efforts, show more favourable performance trade-off. A further value to measure is the disturbance-recovery time. The minimax controller is able to return to the desired operating region faster than all the other approaches. This shows the benefit of making the use of disturbance uncertainty as a part of the optimization problem practical as summarized in Table 3.

Table 3. Comparative Assessment

Criterion	**Conventional Feedback**	**Nominal Optimal**	**Proposed Minimax**
Tracking accuracy	Moderate	Good	Very Good
Disturbance rejection	Low	Moderate	High
Control effort	High	Moderate	Low
Recovery speed	Slow	Moderate	Fast
Robustness	Limited	Moderate	High
Patient adaptation	Limited	Limited	Included
Safety constraints	Basic	Basic	Explicitly incorporated

The minimax formulation is systematic and balances the tracking accuracy, therapeutic action and robustness. Additionally, the feedback law is adaptively adjusted to patient-specific dynamics by adding in patient-specific parameters. But the conclusions from the numbers should be considered only from a mathematical point of view, not from a clinical point of view, when used to validate the proposed control framework. To deploy clinically, there are physiological datasets, patient models with clinical validation, safety certification and prospective clinical evaluation to be made.

8. Conclusion

A mathematical minimax differential-game formulation for personalized closed-loop therapeutic control in smart healthcare systems has been developed in this paper. The patient was thought of as a dynamic system, in which the physiological condition changes according to therapy, the patient's physiological parameters and unknown disturbances. The therapeutic controller was described as the player who has to be minimised, and the disturbance was described as the player who was to be maximised in a zero-sum differential game. While disturbance attenuation was also taken into account, a quadratic cost functional was developed to simultaneously penalise the error due to physiological tracking and also the excessive effort necessary for therapy. It yielded the minimax value function, which was used to get the optimal therapeutic control and worst-case disturbance strategies from the Hamilton-Jacobi-Isaacs (HJI) formulation.

The solution of the HJI equation for the linear-quadratic case was reduced to a Riccati differential equation, which was computationally tractable. Conditions of existence and uniqueness of the closed-loop trajectory, asymptotic stability, boundedness and robustness in the presence of finite-energy disturbances were established by the mathematical analysis. Patient-specific parameter adaptation was then implemented for the controller to be able to change its feedback policy when tuned to estimated physiological parameters. Constraints were also added to guarantee that all physiological states and therapeutic interventions are kept within predefined regions of values.

The numerical analysis confirmed the benefits of using the minimax strategy over conventional and nominal optimal control strategies and especially when facing an unfavourable disturbance. The representative simulation results show that the proposed framework helps to enhance the tracking performance and speed up the disturbance recovery with decreasing control effort. The main novelty of the paper is therefore three-fold, consisting of a combination of minimax optimisation, patient-specific adaptation and safety-constrained closed-loop control in a single mathematical framework. The setup opens the door for additional extensions such as nonlinear models of the patient, real-time estimation of patient states using physiological sensors, time-varying treatment constraints and data-driven parameter identification.

In the future, the mathematically derived results should be confirmed by experimental results and clinical data from patients. More studies could be done on numerical solvers of HJI for high-dimensional nonlinear systems, and the extension to multi-stage therapeutic decision-making. Such advances would further advance the progress from differential-game analysis, a theoretical technique, to clinically deployable personalised closed-loop therapeutic systems.

References

1. N. J. C. Greenwood and J. E. Gunton, “A computational proof of concept of a machine-intelligent artificial pancreas using Lyapunov stability and differential game theory,” *Journal of Diabetes Science and Technology*, vol. 8, no. 4, pp. 791–806, 2014, doi: 10.1177/1932296814536271.

2. T. Başar and P. Bernhard, *H∞-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, 2nd ed. Boston, MA, USA: Birkhäuser, 2008.

3. T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, 2nd ed. Philadelphia, PA, USA: SIAM, 1999.

4. S. Bansal, M. Chen, S. L. Herbert, and C. J. Tomlin, “Hamilton-Jacobi reachability: A brief overview and recent advances,” in *Proc. IEEE Conference on Decision and Control (CDC)*, pp. 2242–2253, 2017, doi: 10.1109/CDC.2017.8263977.

5. M. I. Gomoyunov, “Differential games for fractional-order systems: Hamilton–Jacobi–Bellman–Isaacs equation and optimal feedback strategies,” *Mathematics*, vol. 9, no. 14, p. 1667, 2021.

6. J. Moon, “State and control path-dependent stochastic zero-sum differential games: Viscosity solutions of path-dependent Hamilton–Jacobi–Isaacs equations,” *Mathematics*, vol. 10, no. 10, p. 1766, 2022.

7. B. El Asri and H. Lalioui, “Continuous and impulse controls differential game in finite horizon with Nash-equilibrium and application,” *Journal of Computational and Applied Mathematics*, vol. 424, p. 115009, 2023.

8. M. Sarbaz and W. Sun, “Min–max adaptive dynamic programming for zero-sum differential games,” *International Journal of Control*, vol. 97, no. 12, pp. 2886–2895, 2024, doi: 10.1080/00207179.2024.2309194.

9. A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control barrier functions: Theory and applications,” in *Proc. IEEE European Control Conference (ECC)*, pp. 3420–3431, 2019, doi: 10.23919/ECC.2019.8796030.

10. J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, “A general safety framework for learning-based control in uncertain robotic systems,” *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2019.

11. L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, “Learning-based model predictive control: Toward safe learning in control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 269–296, 2020, doi: 10.1146/annurev-control-090419-075625.

12. S. Bansal and C. J. Tomlin, “DeepReach: A deep learning approach to high-dimensional reachability,” 2020.

13. G. Chen, Y. Wang, and Q. Zhou, “Optimal H∞ control based on stable manifold of discounted Hamilton-Jacobi-Isaacs equation,” 2024.

14. X. Zhang and Q. Wang, “Zero-sum differential games guidance law accounting for impact-angle-constrained using adaptive dynamic programming,” *Journal of Intelligent & Robotic Systems*, 2025.

15. L. Zhang, M. Ghimire, Z. Xu, W. Zhang, and Y. Ren, “Pontryagin neural operator for solving general-sum differential games with parametric state constraints,” *Proceedings of the 6th Annual Learning for Dynamics & Control Conference*, vol. 242, pp. 1728–1740, 2024.

16. R.Indhumathi, Ali Bostani, Aravindan Srinivasan, R.Sindoori, Ilyos Abdullayev, Sirojiddin Abrorov, & Mamaev Gulom. (2025). Hybrid Quantum–Cloud Framework for Nonlinear Clustering Optimization and Intelligent Resource Management. *Results in Nonlinear Analysis*, *8*(3), 26–35.

17. S. A. Brown *et al.*, “Six-month randomized, multicenter trial of closed-loop control in type 1 diabetes,” *New England Journal of Medicine*, vol. 381, no. 18, pp. 1707–1717, 2019.

18. M. D. Breton *et al.*, “A randomized trial of closed-loop control in children with type 1 diabetes,” *New England Journal of Medicine*, vol. 383, no. 9, pp. 836–845, 2020.

19. T. Zhu, K. Li, P. Herrero, and P. Georgiou, “Basal glucose control in type 1 diabetes using deep reinforcement learning: An in silico validation,” 2020.

20. L. Ortmann, D. Shi, E. Dassau, F. J. Doyle III, B. J. E. Misgeld, and S. Leonhardt, “Automated insulin delivery for type 1 diabetes mellitus patients using Gaussian process-based model predictive control,” 2021.

21. G. Wang *et al.*, “Optimized glycemic control of type 2 diabetes with reinforcement learning: A proof-of-concept trial,” *Nature Medicine*, vol. 29, pp. 2633–2642, 2023, doi: 10.1038/s41591-023-02552-9.