Epidemics control model with consideration of seven-segment population model

Ali Nasir¹ · Huma Rehman Baig¹ · Muhammad Rafiq¹

Received: 30 April 2020 / Accepted: 6 September 2020 / Published online: 11 September 2020
© Springer Nature Switzerland AG 2020

Abstract

This paper discusses a novel method for modeling the spread of an epidemic that facilitates the calculation of the optimal control policy. The proposed model considers seven compartments in the population as opposed to popular approaches based on three or four compartments. The usual compartments, i.e., susceptible, exposed, infected, and recovered individuals have been included in the seven compartment model with the addition of compartments pertaining to individuals under treatment, vaccinated individuals, and individuals in quarantine. The addition of new compartments allows for the incorporation of multiple control options such as treatment, quarantine, and vaccination. The mathematical expressions involved in the model have been described followed by a discussion on the calculation of optimal control policy. Finally, a simulation-based example has been included for demonstrating the effectiveness of the control policy resulting from the proposed model.

Keywords Markov decision process · Epidemiology · Optimal control

1 Introduction

Control of epidemic infections has been an active area of research in the past two decades or so [1] because it has the potential of saving lives with the minimal possible expense and timely efforts. Many control techniques have been developed for control in epidemiology [1, 2]. For example, predictive (or model predictive control based) techniques [3, 4], stochastic techniques [5, 6], Markov chain based techniques [7, 8], Markov decision processes (MDP) based techniques [9–11], optimal control (mostly Pontryagin's maximum principle) based techniques [2, 12], and recently, network-based techniques [13–15]. This paper proposes an MDP based model. *To the best of our knowledge, this paper is the first to propose a seven compartment MDP model.*

Why would we propose a seven compartment model? The answer to this key question lies in the examination of existing models that are mostly three-compartment (or

sometimes four-compartment), i.e., susceptible, infected, and recovered (for diseases where a recovered individual has immunity to the disease). In such models, the decision of vaccinating or treating an individual is not informed about how many people are already under treatment or have been vaccinated. This means that the size of the population is no longer constant as far as the model is concerned. However, in all the MDP based approaches that we have seen, the size of the population is considered to be constant other than the death or birth rate. Consider (for the sake of understanding) a simple example with a population size of seven individuals. Suppose one of them is infected, five are susceptible, and one is recovered. Now, using a SIR (Susceptible-Infected-Recovered) model, suppose the optimal decision is to vaccinate two of the susceptible individuals. The problem after this decision is that these two (vaccinated individuals) are no longer represented in the model (and they are not dead). At best, in such a case, vaccinated can be counted as recovered

✉ Ali Nasir, a.nasir@ucp.edu.pk | ¹University of Central Punjab, Lahore, Punjab, Pakistan.

by the model but vaccinations may not work 100% of the time so the model will be less realistic (a similar argument can be made for individuals under treatment). Therefore, we have proposed a seven compartment model that is more realistic and represents vaccinated individuals, individuals under treatment, and individuals in quarantine as well as the SIR individuals.

A side effect of making the model more realistic is the added computational complexity. A complete (or thorough) solution to the complexity problem is an avenue of future research. However, we have sketched a possible solution in Sect. 3. Also, once a problem is formulated as an MDP, it is entitled to many approximate dynamic programming algorithms [16–18] for calculation of a quasi-optimal control policy with reduced computational complexity.

The rest of the paper has been organized as follows. Section 2 covers the review of recent and most cited literature relevant to this paper. Section 3 presents the major contribution of this paper, i.e., the MDP-based stochastic model. This section also discusses the computational complexity involved in the proposed model as well as a method (value iteration) of calculating the optimal control policy. Section 4 includes a numerical case study that demonstrates the behavior of an optimal control policy calculated using a value iterations algorithm. Section 5 provides insights regarding the results from the case study and their comparison with the existing models. Section 6 concludes the paper by summarizing the key findings and future research directions.

2 Literature review

There is a vast majority of models in epidemiology where differences between the models is trivial in many cases [19]. The most common model in epidemiology is the SIR (Susceptible-Infected-Recovered) model in which the population is divided into three categories or compartments.

Inherently, the problem of control in epidemiology is a discrete and stochastic one. Therefore, MDP is a suitable candidate for modeling this problem. MDP based models have an added advantage of enabling the calculation of a control policy as opposed to a generic control strategy that does not take into account the current state of the population. For example, the geometric programming based approach in [20] provides a resource allocation strategy based on the maximization of an eigenvalue. But this strategy is generalized and does not take into account the current situation or state.

Continuous-time state space-based models usually employ Pontryagin's maximum principle for calculation of the solution [2, 12]. Differences in such approaches lie in the definition of the cost function or some cases in the

type of disease studied. Population is usually normalized (maximum size = 1) and the objective is to minimize infection at the lowest cost of vaccination or treatment.

Predictive approaches and graph-based approaches have gained popularity in recent years [20, 21]. For example, [22] proposes a virus spread minimization algorithm with some theoretical guarantees based on discrete-time LaSalle invariance principle. Graph theory has also been employed by [13, 14] for distinguishing infections and developing curing policy respectively. Node removal in graphs has been used as a mode of actuation in [15].

There are a few MDP and dynamics Bayesian network-based approaches available in the literature. For example [23] provides an approximate solution for graph-based MDPs but their approach is generic and not focus on the problem of control in epidemiology. Similarly, MDP for breast and ovarian cancer has been proposed in [10]. There has been some work on the stochastic prediction of the epidemic curve [24] as well as on the analysis of epidemics behavior under stochastic perturbations [6], delay [5], and geographical data [25]. A dynamic Bayesian network-based approach is proposed in [26] but this approach is specific to the prognosis of coronary heart disease. The game theory-based approach has been discussed in [27] in the context of the SIRV (Susceptible-Infected-Recovered-Vaccinated) model. MDP based three-compartment model has been discussed recently in [9]. A good reference on stochastic modeling and estimation for epidemics is [28]. An optimized nonlinear control based for SIR models is presented in [29] but the uncertainties involved in the process are not considered therein.

None of the previous MDP models discuss the seven-compartment approach. As discussed in the previous section, such an approach is more realistic and facilitates the control policy in making more informed decisions.

3 Problem formulation and MDP based stochastic model

In this section, we present our proposed MDP model for the following control problem.

Given the statistics and current status of various individuals in a population as infected, susceptible, recovered; device optimal control policy using vaccination, treatment, and isolation as means of preventing the spread of an epidemic disease.

Based on the above problem statement, the main objectives of the research are as follows

- Reduce the number of already infected individuals with the help of treatment
- Control the spread of the disease by using isolation

- Minimize the chances of disease spread using vaccination
- Keep the cost of treatment, vaccination, and isolation at a minimum possible value while meeting the above-mentioned objectives.

3.1 States

The state includes information that is available for decision making. For example, if the information available includes several infected individuals (I), number of individuals susceptible to the disease (S), number of vaccinated individuals (V), number of individuals under treatment (T), number of individuals exposed to the disease but not yet infected (E), number of individuals in quarantine (Q), and number of individuals recovered from the disease (R). The resulting state space becomes,

$$\begin{aligned} X &= \{x_1, x_2, \dots, x_n\}, \\ x_i &= \{S_i, I_i, R_i, V_i, T_i, E_i, Q_i\}, i \in \{1, 2, \dots, n\}, \\ S_i &\in \{0, 1, \dots, N\}, I_i \in \{0, 1, \dots, N\}, R_i \in \{0, 1, \dots, N\}, \\ V_i &\in \{0, 1, \dots, N\}, T_i \in \{0, 1, \dots, N\}, E_i \in \{0, 1, \dots, N\}, \\ Q_i &\in \{0, 1, \dots, N\}, S_i + I_i + R_i + V_i + T_i + E_i + Q_i = N, \forall i. \end{aligned} \quad (1)$$

Now, the number of states $n = \frac{(N+6)!}{N!6!}$ (it is a solved problem in combinatorics that number of possible ways of placing N unlabeled objects in m labeled baskets is $\binom{N+m-1}{N}$). Each state is now a tuple of seven variables. However, to determine the state, it is enough to know the values of any six out of the seven variables in the tuple. One question that arises here is that of computational complexity and scalability of the representation concerning the population size.

3.1.1 Discussion on scalability

Whenever a problem is formulated as an MDP, the size of the state space has to be kept in check because of the curse of dimensionality involved in the methods used to calculate optimal control policy. Therefore, it is important to discuss the scalability of the proposed modeling. Figure 1 shows how the size of the state space varies as a function of population size. In this figure, m is the number of compartments (e.g. S, I, R, V , etc.) in the state space, N is the population size, and n is the size of the state space. The size of the state space is plotted in logarithm scale for clarity of representation. In the experience of the authors, a core i5 laptop with six gigabytes of random access memory can handle up to ten million states. This means that for a three-compartment model, we can handle a population size of about 3500. Whereas for five and

Fig. 1 Size of the state space versus population size

seven compartment models, the limit on N is 120 and 40 respectively.

Population in real cities is of the order of several million. Even in small towns, the population is of the order of one hundred thousand. So the question arises that how an MDP model would handle such large numbers. A simple approach to this situation is state abstraction. Specifically, one unit of population in the state space may represent more than one person. In this way, a population size of 105,000 can be represented with 3500 units of population where each unit represents 30 persons. A drawback of such abstraction is that if any number of persons between one and thirty get infected, the unit will show either zero or thirty people being infected. This is similar to a well-known term quantization error in digital systems. The worst-case error is 15 persons for the case where each unit of the population represents 30 persons. Now if we represent this worst-case error as a percentage of the total population, it is merely 0.014% which should be acceptable for policy calculations. Handling models with five or seven compartments is still not easy. Extending the same example with a population of 105,000 to a seven compartment model with 40 units (instead of 3500), we get an error of 2.5% (as opposed to 0.014% in three-compartment cases). Further study can be done to explore ways of reducing effective computational complexity involved in the problem without incorporating prohibitive error. In this regard, the techniques in Approximate Dynamic Programming (ADP) [16, 30] may prove to be helpful. Specifically, [30] introduces a decomposition based Approximate Dynamic Programming approach that has been applied to a spacecraft control problem but can be adapted for epidemics control problem discussed in this paper. The key is to decompose the population into appropriate subgroups.

The diagram shows a directed graph with two rows of nodes. The top row contains nodes S, I, R, V, E, T, Q and the bottom row contains nodes $S', I', R', V', E', T', Q'$. Solid arrows represent direct mappings: $S \rightarrow S', I', R', E', T'$; $I \rightarrow I', R', V', E', Q'$; $R \rightarrow R'$; $V \rightarrow V'$; $E \rightarrow E'$; $T \rightarrow T'$; and $Q \rightarrow Q'$. Dashed arrows represent indirect mappings: $V \rightarrow I'$ and $V \rightarrow E'$.

3.4 Reward/cost function

First, we discuss the elements in the problem that have a positive or negative impact on the solution. The most important element in the cost is the number of infected individuals (I). Secondly, when the vaccination is available, the number of susceptible individuals should also be part of the cost function otherwise the optimal control would never use vaccination [9]. The purpose of vaccination is to build immunity among the susceptible individuals and once the immunity is created, a susceptible individual (S) is labeled as recovered (R) in the model. This is because the presence of immunity (or antibodies) against the disease is the only difference between the susceptible and recovered individuals. Having more recovered individuals is desirable because it reduces the chances of the spread of the disease. Therefore, it is important to have a cost function term associated with the susceptible individuals so that the optimal policy can justify spending money on vaccination. Next, keeping individuals in quarantine is costly and hence Q should be in the cost function. The cost of treatment and vaccination can be included by using either variable T and V or the cost can be included in the actions u_v and u_t for vaccination and treatment respectively. Finally, we must include the cost of exposed individuals (E) in the cost function since exposure leads to infection. The resulting cost function can be written as

$$C(x, u) = \alpha_0 I + \alpha_1 S + \alpha_2 E + \alpha_3 T + \alpha_4 Q + \beta_k u_{vk}. \quad (5)$$

where α_i ($i \in \{0, 1, 2, 3, 4\}$) and β_k ($k \in \{1, 2, \dots, N\}$) are positive constants. Also u_{vk} is k th vaccination action as presented in (2). Notice that the above cost function is an incremental cost. Total cost incorporated in an MDP problem is the expected sum of the cost over the whole decision horizon given as

$$C^* = \min_{u^h \in U} \left(E \left[\sum_{h=0}^{\infty} \gamma^h C(x^h, u^h) \right] \right). \quad (6)$$

In the above equation, h is the decision horizon i.e. the number of decisions to be made before the end of the problem. In (6), h varies from 0 to infinity indicating an infinite horizon. In general, the horizon can be finite rather in some cases, the horizon should be finite. Such cases arise in the problems where the deadline for decision making is predefined or the number of decisions allowed is limited such as in blackjack (cards game).

The cost function defined in (5) is a linear function of the variables in state and actions. In general, the cost does not have to be linear. A comparison of two different types of reward functions is discussed in [9]. For example, the cost of having patients may depend upon

thresholds such as for the first few patients, the cost is low, then the number of patients above a certain threshold cost more and so on.

Furthermore, one could use exponential or quadratic functions instead of linear to reflect any real effect of the problem specific to the time and place (town/city) for which the optimal policy is to be calculated.

3.5 Calculation of the optimal control policy

In this section, we discuss the value iterations method for the calculation of optimal policy using an MDP model.

As the name suggests, value iteration is an iterative technique that is based on the value function of the states. The resulting control policy is optimal concerning the expected value of the criteria provided in (6). The decision making horizon h in value iteration is infinite. Value of each state is given by

$$V(x) = \max_u \left(-C(x, u) + \gamma \sum_{x' \in X} P(x'|x, u) V(x') \right) \quad (7)$$

The value of a state depends upon the sum of the cost incurred by that state and the value of the states that it can lead to. This definition of the value function converges to the optimal value for each state via iterative calculations of the form

$$V_{h+1}(x) \leftarrow \max_u \left(-C(x, u) + \gamma \sum_{x' \in X} P(x'|x, u) V_h(x') \right) \quad (8)$$

Here γ is the discount factor that ranges between zero and one. In case of epidemic infections, taking quick actions is desirable hence the discount factor value should be low e.g. between 0.5 and 0.8. Discount factor with value 0.5 means that the value of reaching the state decreases by half in each passing decision. A low discount factor results in less iteration before the convergence is achieved.

Once the optimal value is obtained, the optimal policy is calculated using the following expression

$$\Pi^*(x) = \arg \max_u \left(-C(x, u) + \gamma \sum_{x' \in X} P(x'|x, u) V^*(x') \right) \quad (9)$$

Here $V^*(x')$ represents the optimal value of state x' and $\Pi^*(x)$ represents an optimal policy for state x . Note that the optimal policy, in this case, is stationary, i.e., the optimal decision at each state is independent of the time at which the state is reached. This is because, in infinite decision horizon, time is irrelevant. The only factor affecting the early verses late decisions is the discount factor.

Table 1 Parameter values for the case study

Parameter Name	Value
N (size of scaled population)	25
x (total number of states)	736,281
p (number of partitions in actions)	1 (each action applied to a single individual of the scaled population)
α_0 (weighting factor in (5))	100
β (weighting factor in (5))	20
α_1 (weighting factor in (5))	40
α_2 (weighting factor in (5))	70
α_3 (weighting factor in (5))	40 for $T \leq 10$, 50 for $10 < T \leq 20$, and 60 for $T > 20$
α_4 (weighting factor in (21))	45

4 Case study

In this section, shows how the optimal policy can be summarized and analyzed parametrically. Also, the sample trajectory has been shown to demonstrate the performance of the optimal policy in terms of confining and diminishing the epidemic disease. The case study presented here involves 736,281 states where each state has four possible actions and each action can lead to one of the 384 possible next states (all possible combinations of the transitions as in Fig. 2). The values of the parameters used in the simulations are summarized in Table 1. Note that the values presented in the table are very important as the behavior of the optimal policy is directly related to these values. It is also important to understand that it is not the exact value that matters (as far as the optimal policy is concerned) rather, the key factor is the *relative value* of the parameters. For example, the cost of having an infected individual ($\alpha_0 = 100$) is five times the cost of vaccination ($\beta = 20$). This provides a perspective of how important it is to avoid (or reduce) the infection among the individuals compared to using vaccination resources. Similarly, the cost ($\alpha_1 = 40$) of having a susceptible (S) is less than that ($\alpha_2 = 70$) of having an exposed individual (E) indicating that an exposed individual is relatively more undesirable compared to the susceptible one. A special piece-wise continuous value assignment has been associated with α_3 (the cost of treatment) to emphasize that as the number of patients to be treated goes up, so does the cost of treatment. This assignment is realistic because as the demand for the supplies in the hospitals increases, the cost increases as well. Overall, the values in Table 1 have been selected to present a realistic approach. However, for specific diseases, the values may be estimated based on the actual cost of particular vaccination for the disease. The cost of hospitalization may also vary from country to country and even within a country, it may vary depending upon the type and standards of a particular hospital.

The values related to the transition probabilities are presented in Table 2. The first column of the table shows the label transition (please refer to Fig. 2) for a single unit of population. The remaining four columns represent the corresponding probabilities of these transitions for each of the four possible actions. For example, the probability that a susceptible unit of population shall stay susceptible is 0.65 if no action (*NOOP*) is executed. Similarly, the probability of recovery of an infected individual is 0.1 for *NOOP*. The sum of all possible transitions from each of the seven variables is one. The probability values presented in this section have been inspired by the flu statistics¹ in the United States of America (USA). For example, about 20% of the population (in worst case) is infected by flu in the USA (the probability value of 0.25 in the second row of Table 2). Some values have been exaggerated a little bit to indicate that the disease is not easily recovered. The reason behind the selection of pessimistic probability values is to put the resulting optimal policy to a tough test.

The MDP for the seven-compartment model is solved using value iteration. The resulting optimal policy is summarized in Table 3. This table summarizes which actions are deemed optimal more often than the others i.e. it can be seen that vaccination and treatment are deemed optimal quite often.

A deeper parametric analysis of the policy is presented in Table 4. The values (or ranges) in this table present insights into the behavior of the optimal policy. For example, if the number of susceptible units of population (S) is greater than five, then it is not optimal to do nothing. Similarly, *NOOP* is not an optimal choice if the number of infected units in the population (I) is greater than two. Another interesting insight is that vaccination is no longer optimal when " I " is greater than ten. This means that when there are quite a few people infected, treating the infected individuals takes precedence over vaccinating

¹ <https://www.webmd.com/cold-and-flu/flu-statistics>.

Table 2 Transition Probabilities for a unit of population given against possible actions

State Transition	Probability value (NOOP)	Probability value (vaccination)	Probability Value (treatment)	Probability value (quarantine)
$S \rightarrow S$	0.65	0.05	0.65	0.65
$S \rightarrow I$	0.25	0	0.25	0.25
$S \rightarrow V$	0	0.95	0	0
$S \rightarrow E$	0.1	0	0.1	0.1
$I \rightarrow I$	0.9	0.9	0	0
$I \rightarrow R$	0.1	0.1	0	0
$I \rightarrow T$	0	0	1	0
$I \rightarrow Q$	0	0	0	1
$R \rightarrow R$	1	1	1	1
$V \rightarrow I$	0	0	0	0
$V \rightarrow V$	1	1	1	1
$V \rightarrow E$	0	0	0	0
$E \rightarrow I$	0.5	0.5	0.5	0.5
$E \rightarrow E$	0.5	0.5	0.5	0.5
$T \rightarrow R$	0.3	0.3	0.3	0.3
$T \rightarrow T$	0.7	0.7	0.7	0.7
$Q \rightarrow R$	0.1	0.1	0.1	0.1
$Q \rightarrow Q$	0.9	0.9	0.9	0.9

Table 3 Summary of the optimal policy

Optimal decision	Number of states
NOOP	44,812
Vaccination	2,43,263
Treatment	4,45,314
Quarantine	2892

the susceptible ones. In our case study, treatment is highly preferred over quarantine. This is because treatment yields a higher probability of recovering. In another setting where treatment is too costly or the spread probability of the infection is too high, quarantine may be preferable over treatment.

The traditional way of analyzing the control policy for epidemic infections is to show how it renders the susceptible and infected individuals to zero asymptotically. Therefore, a similar analysis of the proposed approach has also been presented. Figures 3 and 4 show the results of the optimal control policy with initially 20 units of population

being susceptible and five units being infected. Exogenous events in these simulations have been selected using uniform distribution (an exogenous event refers to a possible state transition in the model given current state and action). Notice that the infected individuals recover after treatment and quarantine the susceptible units of the population are vaccinated. This kind of insight and those in Table 4 could be used to determine soft rules for epidemic control (the details of the development of such rules are beyond the scope of this paper).

5 Discussion of the results

A major advantage of the proposed approach over the existing approaches is the inclusion of more detailed information in the model. Although the case study has been designed using pessimistic values of recovery rate (see Table 2), still the numerical results (Figs. 3 and 4) indicate good control over the spread of the disease without a huge peak in the number of individuals under

Table 4 Parametric analysis of the optimal policy

Action	Variable range	S	I	R	V	T	E	Q
NOOP		0–5	0–2	0–25	0–25	0–24	0–25	0–24
Vaccination		1–25	0–10	0–24	0–24	0–24	0–24	0–24
Treatment		0–23	0–25	0–24	0–24	0–22	0–23	0–25
Quarantine		0–19	0–5	0–24	0–24	0–25	0–24	0–0

Fig. 3 Results from the optimal policy with initial condition $S=20$, $I=5$

Fig. 4 Results from the optimal policy with initial condition $S=20$, $I=5$

treatment or number of individuals being vaccinated or quarantined at any particular instant. Compared to the proposed approach, the SIR model-based approach [9] cannot depict how many individuals are currently being treated, and therefore, the decision to hospitalize another individual is less informed compared to the proposed approach. As a result, there is a risk that the optimal policy ends up choking the medical support system. Similarly, the other approaches such as [3, 15] do not consider hospitalization and treatment as part of the model and hence the decisions are not well informed. Since the proposed model incorporates the

information of the individuals under treatment, the individuals in quarantine, the exposed individuals, and the associated transition probabilities, the resulting optimal policy successfully avoids large peaks in the number of individuals under treatment or quarantine at any given time. Thus avoiding the overburdening of the medical support system.

The incorporation of additional information in the model comes at the cost of computational complexity. This drawback, however, is subdued by the advancements in Nanotechnology, cloud computing, and available literature on approximate dynamic programming [16]. Another key advantage of the proposed approach is that the optimal policy is calculated once (offline) and the same can be used for any situation (or state) that may arise in real-life scenarios (as long as the state is within the scope of the model). On the other hand, the control approaches discussed in [2] for example provide a control law that is applied in an online manner.

Finally, the insights developed using the optimal policy such as those presented in Tables 3 and 4 can be used to reverse engineer the best values of the weighting parameters in the cost function as it might be difficult in some cases to determine the cost of having an infected individual in the population. Such a provision makes the proposed approach preferable over the differential equation (or difference equation) based approaches in the literature [2, 15, 21]. Specifically, the boundaries (or thresholds) relating possible optimal actions with the number of units of the population in a certain compartment and ranges of the number of units of the population for each possibly optimal action. In traditional difference equation-based models, the population could be in fractions and so could be the control signal. Whereas the MDP based policy gives exact numbers (as positive integers) for control signals and population which is easier to interpret, implement, and understand.

6 Conclusion

The research gap addressed in this paper is the absence of key information in existing epidemics control models. As a result, a seven-segment model has been developed as opposed to the existing three or four-segment models. The reason for proposing a seven-segment model is to enable the resulting control policy in making more informed decisions such as controlling the spread of the epidemic while keeping the burden on the medical support system to an acceptable value. The results from the case study indicate that the desired performance can be achieved from the optimal policy that uses the proposed model. Specifically, the optimal policy can control the epidemic spread while

keeping the number of individuals under treatment of quarantine at any particular instant to an acceptable value.

The advantage of informed decisions and good performance comes at the cost of computational complexity. In this regard, existing approaches for computational complexity reduction can be used and more specific approaches can also be developed as a future work. Other future directions regarding the proposed work involve adapting various theoretical results in the existing models to the proposed one. Also, a multilayer approach to decision making can be researched that uses low level and high-level decision making to facilitate distributed or networked control.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

- Nowzari C, Preciado VM, Pappas GJ (2016) Analysis and control of epidemics: a survey of spreading processes on complex networks. *IEEE Control Syst* 36(1):26–46
- Sharomi O, Malik T (2015) Optimal control in epidemiology. *Ann Oper Res* 251(1–2):55–71
- Watkins NJ, Nowzari C, Pappas GJ (2018) Robust economic model predictive control of continuous-time epidemic processes. [arXiv:1707.00742v5](https://arxiv.org/abs/1707.00742v5)[math.OC]
- Watkins NJ, Nowzari C, Pappas GJ (2017) Inference prediction and control of networked epidemics. In: *Proceedings of IEEE American Control Conference*, pp 5611–5616
- Liu Q, Jiang D, Shi N, Hayat T, Alsaedi A (2016) Asymptotic behaviors of a stochastic delayed SIR epidemic model with nonlinear incidence. *Commun Nonlinear Sci Numer Simul* 40:89–99
- Ji C, Jiang D (2017) The threshold of a non-autonomous SIRS epidemic model with stochastic perturbations. *Math Methods Appl Sci* 40(5):1773–1782
- Gómez S, Arenas A, Borge-Holthoefer J, Meloni S, Moreno Y (2010) Discrete-time Markov chain approach to contact-based disease spreading in complex networks. *Europhys Lett* 89(3):2010
- Ahn HJ, Hassibi B (2014) On the mixing time of the SIS Markov chain model for epidemic spread. In: *Proceedings IEEE conference on decision and control*, pp 6221–6227
- Nasir A, Rehman H (2017) Optimal control for stochastic model of epidemic infections. In: *2017 14th international Bhurban conference on applied sciences and technology (IBCAST)*. IEEE, 2017
- Abdollahian M, Das TK (2015) An MDP model for breast and ovarian cancer intervention strategies for BRCA1/2 mutation carriers. *IEEE J Biomed Health Inform* 19(2):720–727
- Gast N, Gaujal B, Le Boudec J (2012) Mean field for markov decision processes: from discrete to continuous optimization. *IEEE Trans Autom Control* 57(9):2266–2280
- Gubar E, Zhu Q (2013) Optimal control of influenza epidemic model with virus mutations. In: *Proceedings of European control conference*, pp 3125–3130
- Drakopoulos K, Ozdaglar A, Tsitsiklis J (2014) An efficient curing policy for epidemics on graphs. In: *53rd IEEE conference on decision and control*, Los Angeles, CA, 2014, pp 4447–4454
- Milling C, Caramanis C, Mannor S, Shakkottai S (2015) Distinguishing infections on different graph topologies. *IEEE Trans Inf Theory* 61(6):3100–3120
- Watkins NJ, Pappas GJ (2019) Control of generalized discrete-time SIS epidemics via submodular function minimization. *IEEE Control Syst Lett* 3(2):314–319
- Powell WB (2007) *Approximate dynamic programming: solving the curses of dimensionality*, vol 703. Wiley, Hoboken
- Parr R (1998) Flexible decomposition algorithms for weakly coupled Markov decision problems. In: *Proceedings of the fourteenth conference on uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc
- Anders J, Barto A (2005) A causal approach to hierarchical decomposition of factored MDPs. In: *Proceedings of the 22nd international conference on machine learning*. ACM
- Hethcote HW (1994) A thousand and one epidemic models. In: Levin SA (ed) *Frontiers in mathematical biology (series lecture notes in biomathematics, no. 100)*. Springer, New York, pp 504–515
- Preciado VM, Zargham M, Enyioha C, Jadbabaie A, Pappas GJ (2014) Optimal resource allocation for network protection: a geometric programming approach. *IEEE Trans Control Netw Syst* 1(1):99–108
- Nowzari C, Preciado VM, Pappas GJ (2014) Stability analysis of generalized epidemic models over directed networks. In: *Proceedings of conference on decision and control*, Los Angeles, CA, December 2014, pp 6197–6202
- Ramirez-Llanos E, Martinez S (2014) A distributed algorithm for virus spread minimization. In: *Proceedings of American control conference*, Portland, OR, 2014, pp 184–189
- Sabbadin R, Peyrard N, Forsell N (2012) A framework and a mean-field algorithm for the local control of spatial processes. *Int J Approx Reason* 53(1):66–86
- Zamiri A, Yazdi HS, Goli SA (2015) Temporal and spatial monitoring and prediction of epidemic outbreaks. *IEEE J Biomed Health Inform* 19(2):735–744
- Canino G, Guzzi PH, Tradigo G, Zhang A, Veltri P (2017) On the analysis of diseases and their related geographical data. *IEEE J Biomed Health Inf* 21(1):228–237
- Orphanou K, Stassopoulou A, Keravnou E (2016) DBN-extended: a dynamic Bayesian network model extended with temporal abstractions for coronary heart disease prognosis. *IEEE J Biomed Health Inf* 20(3):944–952
- Reluga TC, Galvani AP (2011) A general approach for population games with application to vaccination. *Math Biosci* 230(2):67–78
- Andersson H, Britton T (2012) *Stochastic epidemic models and their statistical analysis*, vol 151. Springer, Berlin
- Guang Y (2010) Synthesize control for an SIR model with non-linear saturation infectious force. In: *2010 Chinese control and decision conference*. IEEE
- Nasir A, Atkins EM, Kolmanovsky I (2017) Robust science-optimal spacecraft control for circular orbit missions. In: *IEEE transactions on systems, man, and cybernetics: systems*, vol 50(03), pp 923–934

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.