







Efficient Evaluation of Influenza Mitigation Strategies Using Preventive Bandits

Pieter Libin^{1,2} , Timothy Verstraeten¹ , Kristof Theys² ,
Diederik M. Roijers¹ , Peter Vrancx¹ , and Ann Nowé¹ 

¹ Artificial Intelligence lab, Department of Computer Science,
Vrije Universiteit Brussel, Brussels, Belgium

{pieter.libin,tiverstr,diederik.roijers,peter.vrancx,ann.nowe}@vub.ac.be

² Department of Microbiology and Immunology, Rega Institute
for Medical Research KU Leuven - University of Leuven, Leuven, Belgium
kristof.theys@kuleuven.be

Abstract. Pandemic influenza has the epidemic potential to kill millions of people. While different preventive measures exist, it remains challenging to implement them in an effective and efficient way. To improve preventive strategies, it is necessary to thoroughly understand their impact on the complex dynamics of influenza epidemics. To this end, epidemiological models provide an essential tool to evaluate such strategies *in silico*. Epidemiological models are frequently used to assist the decision making concerning the mitigation of ongoing epidemics. Therefore, rapidly identifying the most promising preventive strategies is crucial to adequately inform public health officials. To this end, we formulate the evaluation of prevention strategies as a multi-armed bandit problem. Through experiments, we demonstrate that it is possible to identify the optimal strategy using only a limited number of model evaluations, even if there is a large number of preventive strategies to consider.

Keywords: Epidemiological models · Preventive strategies
Pandemic influenza · Multi-armed bandits · Reinforcement learning

1 Introduction

The influenza virus is responsible for the deaths of half of a million people each year [1]. Additionally, seasonal influenza epidemics cause a significant economic burden [2]. While influenza is typically confined to local epidemics, it can cause a pandemic when a novel strain emerges that has the ability to spread rapidly among a susceptible human host population [3]. Pandemic influenza occurs less frequently than seasonal influenza but the outcome with respect to morbidity and mortality can be much more severe, potentially killing millions of people worldwide [4]. Therefore, it is essential to study mitigation policies to control pandemic influenza epidemics.

For influenza, different preventive measures exist: i.a., vaccination, social measures (e.g., school closures) and antiviral drugs. However, the efficiency of these measures greatly depends on their availability, as well as on epidemiological characteristics. Furthermore, governments typically have limited resources to implement such measures. Therefore, it remains challenging to formulate prevention strategies that make effective and efficient use of these preventive measures while putting as little strain on the available resources as possible.

To improve the development of preventive strategies, it is necessary to thoroughly understand the complex dynamics of influenza epidemics. To this end, epidemiological models are commonly used. Such models study the effects of preventive measures *in silico* [5,6].

Epidemiological models are frequently used to assist the decision making concerning the mitigation of ongoing epidemics (not only for influenza, e.g., the H1N1/09 influenza pandemic [7], but also the 2014–2016 Ebola epidemic [8], the 2016 yellow fever outbreak [9], etc.). Therefore, rapidly identifying the most promising preventive strategies is crucial. This however, can be at odds with the accuracy of the models.

There are two main types of epidemiological models that are frequently applied: *compartment models*, which divide the population into discrete homogeneous states (i.e., compartments) and describe the transition rates from one state to another, and *individual-based models* that explicitly represent all individuals and their connections, and simulate the spread of a pathogen among these individuals. While individual-based models are usually associated with a greater model complexity and computational cost than compartment models, they allow for a more accurate evaluation of preventive strategies [10–12]. It is therefore highly preferable to use individual-based models whenever computational resource constraints permit. In order to make it feasible to use individual-based models, it is essential to use the available computational resources as efficiently as possible.

The outcome of the simulation of a preventive strategy in a stochastic individual-based model, is a sample of that strategy’s *outcome distribution*. In the literature, a set of possible prevention strategies is typically evaluated by simulating each of the strategies a predefined number of times (e.g., [13]). However, this can allocate a large proportion of computational resources to explore the effects of highly sub-optimal strategies.

We therefore propose to apply *reinforcement learning* [14] with *multi-armed bandits* [15]. Reinforcement learning is the study of how to balance exploitation (i.e., further simulating the effects of what we believe to be the best preventive strategy to obtain more accurate results) and exploration (i.e., simulating the effects of other strategies to see whether they might actually be better than our current best). By using this framework, we aim to reduce the number of required model evaluations to determine the most promising preventive strategies. This reduces the total time required to study a given set of prevention strategies, making the use of individual-based models attainable in studies where it would otherwise not be computationally feasible. Additionally, faster evaluation can

also free up computational resources in studies that already use individual-based models, capacitating researchers to explore different model scenarios. Considering a wider range of scenarios increases the confidence about the overall utility of prevention strategies.

In this paper, we formulate the evaluation of preventive strategies as a multi-armed bandit learning problem in Sect. 3. The utility of this new method is confirmed through experiments in the context of pandemic influenza in Sect. 4, using the popular FluTE individual-based model [10]. Our results show that we can quickly focus our computational resources on the optimal prevention strategy. We thus conclude that our method has the potential to be used as a decision support tool for mitigating influenza epidemics.

2 Background

This section provides background on the application domain (i.e., finding mitigation strategies for pandemic influenza using epidemiological models) and learning methods (i.e., multi-armed bandits) approached in this study.

2.1 Pandemic Influenza

Influenza is an infectious disease caused by the influenza virus. The primary prevention strategy to mitigate seasonal influenza is to produce vaccine prior to the epidemic, anticipating the virus strains that are expected to circulate. This vaccine pool is used to inoculate the population before the start of the epidemic. As influenza viruses are constantly evolving, the stockpiling of vaccine to prepare for a pandemic is not possible, as the vaccine should be specifically tailored to the virus that is the source of the pandemic [16]. Therefore, before an appropriate vaccine can be developed, the responsible virus needs to be identified [16]. Hence, vaccine will be available only in limited supply at the beginning of the pandemic [16]. Additionally, vaccine shortage can be induced by problems with vaccine production (e.g., the vaccine contamination in the United States in 2004–2005 [17]). While pandemic influenza has been studied and modeled extensively, there are still many aspects with respect to mitigation strategies that remain to be investigated [13, 18]. Furthermore, awareness was raised recently about certain parameters and assumptions used in epidemiological models to be too conservative to explore the full epidemic potential of pandemic influenza, and as a result evaluate mitigation strategies overly optimistic [19]. These concerns indicate that the reevaluation of preventive strategies, taking into account more realistic assumptions, is warranted.

The severity of pandemic influenza, the limited availability of vaccine and an extensive set of open research questions render this field a primary target to evaluate preventive strategies more efficiently.

2.2 Epidemiological Models

Epidemiological models are an indispensable tool to investigate how pathogens spread through a population and to evaluate mitigation strategies. Epidemiological models are therefore crucial tools to assist policy makers with their decisions [20, 21]. Modeling epidemiological processes can be approached by means of individual-based models or compartment models. Compartment models divide the population into discrete homogeneous states (i.e., compartments) and describe the transition rates from one state to another. Compartment models can be formulated as differential equations and thus form a mathematical framework to model epidemics. Individual-based models, on the other hand, explicitly represent all individuals and their connections and simulate the spread of a pathogen among this network of individuals. Individual attributes influence the way the contact network evolves temporally and spatially. Additionally, the infection progress and the different stages associated with this progress is modeled per individual. Individual-based models allow to evaluate therapeutic and preventive interventions on the level of individuals. Compartment models generalize on population level and represent the expectation of epidemiological outcomes, while individual-based models are able to represent individual heterogeneity. Modeling a greater level of heterogeneity is usually associated with a greater model complexity and computational cost, but allows for a more accurate evaluation of preventive strategies [10–12, 22, 23]. The result of a model evaluation is referred to as the model outcome. The relevant model outcomes greatly depend on the policy makers’ research questions (e.g., prevalence, proportion of symptomatic individuals, morbidity, mortality, cost).

2.3 Modeling Influenza

There is a long tradition to use individual-based models to study influenza epidemics [5, 6, 13], since it allows for a more accurate evaluation of preventive strategies. A main example is FluTE [10], an influenza individual-based model that has been the driver for many high impact research efforts over the last decade [5, 6, 24]. FluTE implements a contact model where the population is divided into communities of households [10]. The population is thus organized in a hierarchy of social mixing groups where the contact intensity is inversely proportional with the size of the group (e.g., closer contact between members of a household than between colleagues). FluTE also supports worker’s commute and the travel of individuals, both model components that can be parameterized from census data. FluTE’s contact network can be informed by population census data, and geographical regions as large as the United States can be modeled [10]. Next to the social mixing model, FluTE implements an individual disease progression model, where different disease stages are associated with different levels of infectiousness. To support the evaluation of prevention strategies, FluTE allows the simulation of both therapeutic interventions (i.e., vaccines, antiviral compounds) and non-therapeutic interventions (i.e., school closure, case isolation, household quarantine). FluTE is a highly customizable simulator in which all model components can be configured in great detail.

2.4 Multi-armed Bandit

The multi-armed bandit problem [25] concerns a k -armed bandit (i.e., a slot machine with k levers) where each arm A_i returns a reward r_i when it is pulled. As each arm returns rewards according to a particular reward distribution, a gambler wants to play a sequence of arms to maximize her/his reward. A strategy to play such a sequence of arms is called a *policy*. Such policies need to carefully balance between exploitation (i.e., choose the arms with the highest expected reward) and exploration (i.e., explore the other arms to potentially identify even more promising arms).

Multi-armed bandits have been proven useful to model many empirical cases: i.a., the organization of clinical trials such that patient mortality is minimized [26], resource allocation among competing stakeholders [27], adaptive routing [28], A/B testing [29] and automated auctioning [30].

One of the simplest bandit learning algorithms is the ε -greedy policy [14], this policy selects the greedy arm (i.e., the arm with the highest expected reward) with probability $1 - \varepsilon$ and explores the non-greedy arms with probability ε . Another popular policy is UCB1 (i.e., Upper Confidence Bound) [15]. UCB1 considers the uncertainty of each arms' value (i.e., the uncertainty of the expected reward) by selecting the arm with the highest upper confidence bound. The upper confidence bound for an arm A_i is computed as $\bar{x}_i + \sqrt{c \frac{\ln(n)}{n_i}}$ where \bar{x}_i is the sample average of A_i , n_i is the number of times A_i was played and n is the overall number of plays [15]. The second term is an exploratory term, which decreases when arm A_i is being pulled sufficiently. This promotes the exploration of arms for which the estimated expected reward is uncertain.

3 Methods

To optimize the evaluation of prevention strategies, it is important to identify the best strategy using a minimal amount of model evaluations. Therefore, we propose to formulate the evaluation of prevention strategies as a multi-armed bandit problem. The presented method is generic with respect to the kind of epidemic that is modeled (i.e., pathogen, contact network, preventive strategies). The method is evaluated in the context of pandemic influenza in the next section.

3.1 Preventive Bandits

Definition 1. *A multi-armed bandit problem [15] consists of $n = |\{A_0, \dots, A_n\}|$ arms and a (time-independent) reward distribution $P(r|A_i, \theta_i)$ for each arm, where θ_i are the parameters of the distribution. At each time step, t , an agent (i.e., gambler) chooses and plays an arm A_i , and receives a reward, r_t sampled (independently) from $P(r|A_i, \theta_i)$. The reward distributions' parameters are unknown to the agent.*

The goal in a multi-armed bandit is to optimize the cumulative sum of rewards. In order to do so, it must select arms that exploit its current knowledge about θ_i , i.e., by picking the best arm it has seen so far. However, it must also explore, in order to discover arms that are better. Because the rewards are received stochastically, the agent must never exclude the possibility that its current estimates are wrong.

In our setting, we want to find the optimal preventive strategy from a set of strategies by evaluating the strategies in an epidemiological model.

Definition 2. *A stochastic epidemiological model E is a function $\mathcal{C} \times \mathcal{P} \rightarrow \mathbb{R}$ where: $c \in \mathcal{C}$ is a configuration, $p \in \mathcal{P}$ is a preventive strategy and the codomain \mathbb{R} represents the model outcome distribution.*

Note that a model configuration $c \in \mathcal{C}$ describes the entire model environment. This means both aspects inherent to the model (e.g., FluTE’s mixing model) and options that the modeler can provide (e.g., population statistics, vaccine properties, basic reproduction number).

Our objective is to find the optimal preventive strategy from a set of alternative preventive strategies $\{p_0, \dots, p_n\} \subset \mathcal{P}$ for a particular configuration $c_0 \in \mathcal{C}$ (corresponding to the studied epidemic) of a stochastic epidemiological model. To this end, we define a preventive bandit.

Definition 3. *A preventive bandit has $n = |\{p_0, \dots, p_n\}|$ arms. Playing arm p_i corresponds to evaluating $E(c_0, p_i)$ by running a simulation of the epidemiological model. Evaluating $E(c_0, p_i)$ results in a sample of the model outcome distribution. The reward of p_i is a mapping of this model outcome (i.e., a sample of the model outcome distribution) using a mapping function $\mathbb{R} \rightarrow \mathbb{R}$.¹*

A preventive bandit is thus a multi-armed bandit, in which the arms are preventive strategies, and the reward distribution is implemented by an instance of a stochastic epidemiological model $E(c_0, p_i)$. We note that while the parameters of the reward distribution are in fact known, it is intractable to determine the optimal reward analytically from the stochastic epidemiological model.

Formulating the evaluation of preventive strategies in terms of a bandit problem provides us with a new framework to reason about this task. The goal is to determine the best preventive strategy (i.e., the prevention strategy that mitigates the pandemic best on average) using as little model evaluations as possible.

3.2 Identifying the Optimal Strategy

Our goal is to identify the optimal strategy for a particular configuration $c_0 \in \mathcal{C}$ while thoroughly exploring all preventive strategies. For this purpose, we explore the use of the popular ε -greedy and UCB1 algorithms.

¹ The mapping function allows the model outcome to be represented more conveniently for learning.

4 Experiments

Two experiments were composed and performed in the context of pandemic influenza modeling. More specifically, in these experiments we analyze the mitigation strategy to vaccinate a population when only a limited number of vaccine doses is available (details about this scenario in Sect. 2). The experiments are inspired by the work of Medlock [31].

When the number of vaccine doses is limited, it is imperative to identify an optimal vaccine allocation strategy [31]. In our experiments, we explore the allocation of vaccines over five different age groups: pre-school children, school-age children, young adults, older adults and the elderly.

The experiments share a base model configuration, but differ with respect to a key epidemiological parameter: the basic reproduction number (i.e., R_0). The basic reproduction number represents the number of infections that is, by average, generated by one single infection.

4.1 Influenza Model and Configuration

The epidemiological model used in the experiments is the FluTE stochastic individual-based model (for details please refer to Appendix A). FluTE comes with a set of sample populations, in this experiment we use the sample population that describes a single community consisting of 2000 individuals (for details please refer to Appendix A). At the first day of the simulated epidemic, 10 random individuals are infected (i.e., 10 infections are seeded). The epidemic is simulated for 180 days. During this time no more infections are seeded. Thus, all new infections established during the run time of the simulation, result from the mixing between infectious and susceptible individuals. We assume no pre-existing immunity towards the circulating virus variant. We assume there are 100 vaccine doses to allocate (i.e., vaccine for 5% of the population).

In this experiment, we explore the efficacy of different vaccine allocation strategies. We consider that only one vaccine variant is available in the simulation environment. FluTE allows vaccine efficacy to be configured on 3 levels: efficacy to protect against infection when an individual is susceptible (i.e., VE_{Sus}), efficacy to avoid an infected individual from becoming infectious (i.e., VE_{Inf}) and efficacy to avoid an infected individual from becoming symptomatic (i.e., VE_{Sym}). In our experiment we consider $VE_{Sus} = 0.5$ [32], $VE_{Inf} = 0.5$ [32] and $VE_{Sym} = 0.67$ [7]. The influenza vaccine, as most vaccines, only becomes fully effective after a certain period upon its administration, and the effectiveness increases gradually over this period [33]. In our experiment, we assume the vaccine effectiveness to build up exponentially over a period of 2 weeks [33, 34].

We define two experiments: both experiments use the base model configuration as described above. The two experiments differ with respect to their R_0 (i.e., basic reproduction number) parameter. To evaluate our new method, we select 2 values that are used in many studies: $R_0 = \{1.3, 1.4\}$ [5, 10, 31]. Each experiment thus has its own configuration. With respect to the definition of the

epidemiological model (i.e., $E = \mathcal{C} \times \mathcal{P} \rightarrow \mathbb{R}$), we can express these configurations as $c_{R_0=1.3}$ and $c_{R_0=1.4} \in \mathcal{C}$.

4.2 Formulating Vaccine Allocation Strategies

We consider 5 age groups to which vaccine doses can be allocated: pre-school children (i.e., 0–4 years old), school-age children (i.e., 5–18 years old), young adults (i.e., 19–29 years old), older adults (i.e., 30–64 years old) and the elderly (>65 years old). An allocation scheme can be encoded by a Boolean 5-tuple, where each position in the tuple corresponds to the respective age group. When the value is 1 at a position, this denotes that vaccines should be allocated to the respective age group. When the value is 0 at a position, this denotes that vaccines should not be allocated to the respective age group. When vaccine is to be allocated to a particular age group, this is done proportional to the size of the population that is part of this age group.

Some examples: a preventive strategy where no vaccine should be allocated is encoded as $\langle 0, 0, 0, 0, 0 \rangle$, a preventive strategy where vaccine needs to be allocated uniformly across all age groups is encoded as $\langle 1, 1, 1, 1, 1 \rangle$, a preventive strategy where vaccine needs to be allocated exclusively to children is encoded as $\langle 1, 1, 0, 0, 0 \rangle$.

To decide on the best vaccine allocation strategy, we enumerate all possible combinations of this tuple. Since the tuple consists of a sequence of $\{0, 1\}^*$, the tuple can be encoded as a binary number. This enables us to represent the different allocation strategies by integers (i.e., $\{0, 1, \dots, 31\}$).

With respect to the definition of the epidemiological model (i.e., $E = \mathcal{C} \times \mathcal{P} \rightarrow \mathbb{R}$), this set of 32 strategies is a subset of \mathcal{P} .

4.3 An Influenza Bandit

So far, we defined the model configurations (i.e., $c_{R_0=1.3}$ and $c_{R_0=1.4}$) and the set of preventive strategies (i.e., 32 vaccine allocation strategies) to be evaluated.

Now, let us define the *influenza preventive bandit* B_{Flu} : B_{Flu} has exactly 32 arms (i.e., $\{A_0, \dots, A_{31}\}$). Each arm A_i is associated with the allocation strategy for which the integer encoding is equal to i . To conclude the specification of the influenza bandit B_{Flu} , we describe what happens when an arm A_i of B_{Flu} is played:

1. Invoke FluTE with a model configuration $c_0 \in \mathcal{C}$ and the vaccine allocation strategy $p_i \in \mathcal{P}$ associated with the arm A_i (i.e., this is allocation strategy i , using the strategy’s integer representation).²
2. From FluTE’s output, extract the proportion of the population that experienced a symptomatic infection: $\frac{\# \text{ symptomatic individuals}}{\# \text{ individuals}}$.
3. Return a *reward* = $1 - \frac{\# \text{ symptomatic individuals}}{\# \text{ individuals}}$. Note that the reward denotes the proportion of individuals that did not experience symptomatic infection.

² Note that the configuration is serialized as a text file, for details on the format of this file, refer to Appendix B.

4.4 Outcome Distributions

To perform an initial analysis concerning the outcome distributions of the 32 prevention strategies, all strategies were evaluated 1000 times for both model configurations (i.e., $c_{R_0=1.3}$ and $c_{R_0=1.4} \in \mathcal{C}$). Note that generating thousands of samples (i.e., 2×32000 in this case) would not be computationally feasible when considering a larger population. This analysis is performed to identify the best strategy, such that we can properly validate the results from our learning experiments.

The outcome distributions are visualized in Figs. 1 and 2 for $c_{R_0=1.3}$ and $c_{R_0=1.4}$ respectively. A violin plot is used to plot the density of the outcome




Fig. 1. Violin plot that depicts the density of the outcome distribution for 32 vaccine allocation strategies, considering a model environment with $R_0 = 1.3$. For each density, the sample mean is visualized with a diamond. The sample mean of the optimal strategy is depicted with a horizontal line.

distribution per vaccine allocation strategy. The density for a particular strategy is computed based on 1000 samples of the strategy’s outcome distribution. Note that while the distributions have considerable density around the mean of the distribution, there is also quite some density where the outcome is close to 0. This is an artefact of the stochastic simulation: the pathogen is not able to establish an epidemic for certain simulation runs.

Our analysis shows that the best vaccine allocation strategy was identified to be $(0, 1, 0, 0, 0)$ (i.e., vaccine allocation strategy 8) for both model configurations $c_{R_0=1.3}$ and $c_{R_o=1.4}$.




Fig. 2. Violin plot that depicts the density of the outcome distribution for 32 vaccine allocation strategies, considering a model environment with $R_0 = 1.4$. For each density, the sample mean is visualized with a diamond. The sample mean of the optimal strategy is depicted with a horizontal line.

4.5 UCB1 and ε -greedy Experiment

To explore the utility of bandits to evaluate preventive strategies, we average over 500 independent bandit runs for both experiments. For each experiment, we run the ε -greedy ($\varepsilon = 0.1$) and UCB1 algorithm for 1000 iterations.³

The average reward reported in the first experiment is visualized in Fig. 3 for both the ε -greedy and UCB1 algorithm. The average reward reported in the second experiment is visualized in Fig. 4 for both the ε -greedy and UCB1 algorithm.

We observe that the average reward starts to increase from iteration 400, for both ε -greedy and UCB1, and continues to increase for the rest of the iterations. However, we also note that the average reward learning curve increases faster for ε -greedy than for UCB1.




Fig. 3. Reward learning curve for the first experiment (i.e., model with $R_0 = 1.3$), averaged over 500 independent bandits for 1000 iterations. This plot depicts the learning curve for both the ε -greedy and UCB1 algorithms.

In the previous section, the best vaccine allocation strategy was identified to be $\langle 0, 1, 0, 0, 0 \rangle$ (i.e., vaccine allocation strategy 8) for both $c_{R_0=1.3}$ and $c_{R_0=1.4}$. Figure 5 visualizes the percentage of plays of the optimal arm (i.e., vaccine allocation strategy $\langle 0, 1, 0, 0, 0 \rangle$) for the first experiment. Figure 6 visualizes the percentage of plays of the optimal arm (i.e., vaccine allocation strategy $\langle 0, 1, 0, 0, 0 \rangle$) for the second experiment.

³ To remind the reader, each arm involves the invocation of the FluTE simulator, and is therefore associated with a significant computational cost (for details, please see Appendix D).




Fig. 4. Reward learning curve for the first experiment (i.e., model with $R_0 = 1.4$), averaged over 500 independent bandits for 1000 iterations. This plot depicts the learning curve for both the ϵ -greedy and UCB1 algorithms.

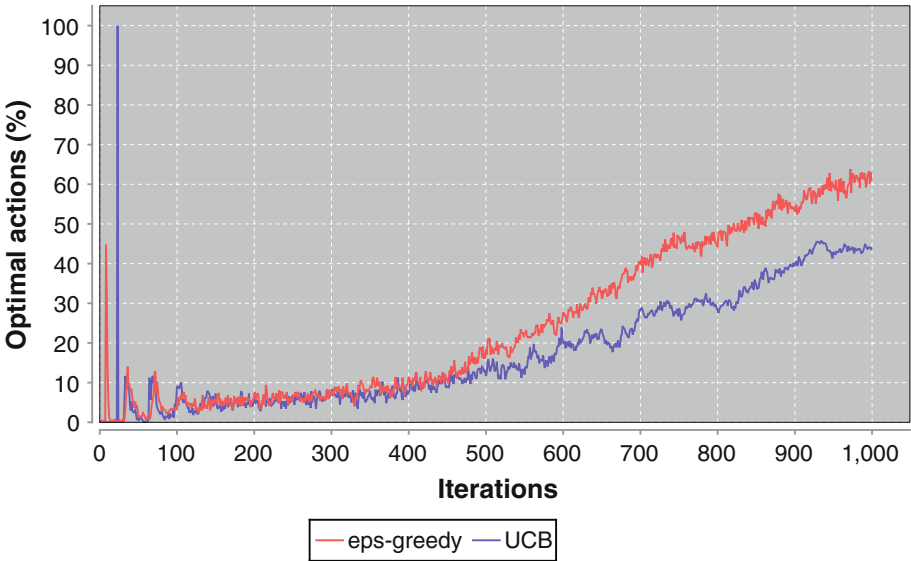


Fig. 5. Optimal action selection learning curve for the first experiment (i.e., model with $R_0 = 1.3$), averaged over 500 independent bandits for 1000 iterations (i.e., the Y-axis depicts the % the optimal action was selected). This plot depicts the learning curve for both the ϵ -greedy and UCB1 algorithms.




Fig. 6. Optimal action selection learning curve for the first experiment (i.e., model with $R_0 = 1.4$), averaged over 500 independent bandits for 1000 iterations (i.e., the Y-axis depicts the % the optimal action was selected). This plot depicts the learning curve for both the ϵ -greedy and UCB1 algorithms.

For both of the experiments, ϵ -greedy ends up selecting optimal actions 60% of the time after 1000 iterations. As we observed for the average reward learning curve, UCB1 also performs worse with respect to the optimal action selection learning curve, reaching only 40–45% optimal action selection.

5 Discussion

Our influenza model, and more specifically the context in which only a limited set of vaccine doses is available, was inspired by the work presented by Medlock and Galvani [31]. However, we consider a much smaller population (i.e., 2000 individuals versus the entire United States), to make it computationally feasible to validate our learning experiments. Furthermore, because of the differences between the model setup presented by Medlock and FluTE, a perfect mapping was not possible. It would therefore not be sound to compare our results directly to the results obtained by Medlock. We were, however, able to reproduce some significant trends. The best strategy identified by our analyses is associated with the allocation of vaccine to children: this is in agreement with Medlock’s work.

The analysis of the outcome distributions for the different vaccine allocation strategies shows that there is one optimal strategy $\langle 0, 1, 0, 0, 0 \rangle$. The differences between the means and medians of the different strategies are however not very pronounced. This is related to the limited number of available vaccine doses.

For both of the experiments, ϵ -greedy ends up selecting optimal actions 60% of the time after 1000 iterations. These results demonstrate that it is possible to identify the optimal strategy using only a limited number of model evaluations, even if there is a large number of preventive strategies to consider. We also observe, that both the average reward and optimal action selection learning curves continue to increase, indicating that the learning has not yet converged. It is however important to stress that, our main interest is not convergence, but to identify the best strategy using a minimal number of model evaluations.

We observe that, in our experiment setting, ϵ -greedy outperforms UCB1, both with respect to the average reward learning curve and the optimal action selection learning curve.

To support the reproducibility of our research, all source code and configuration files used in our experiments is publicly available (for details, please see the Appendices).

6 Conclusions

We formally defined the evaluation of prevention strategies as a multi-armed bandit problem. We used this formal definition to describe a bandit that can be used to evaluate vaccine allocation strategies with the intention to mitigate pandemic influenza. Two elaborate experiments were set up to evaluate this preventive bandit using the popular FluTE individual-based model. To assess the performance of the preventive bandit, we report an average over 500 independent bandit runs, for the two experiments.

We demonstrate that it is possible to identify the optimal strategy using only a limited number of model evaluations, even if there is a large number of preventive strategies to consider.

We are confident that our method has the potential to be used as a decision support tool for mitigating influenza epidemics. To increase this potential, we aim to significantly extend the features of our tool and framework.

Firstly, while our method is evaluated in the context of pandemic influenza, it is important to stress that both our formalisms and infrastructure can be used to evaluate prevention strategies for other infectious diseases. We expect that epidemics of arboviruses (i.e., viruses that are transmitted by a mosquito vector; e.g., Zika virus, Dengue virus) are a particularly interesting use case for our preventive bandits. Only since recently, Dengue and Zika vaccines are available [35] or in the pipeline [36], and the optimal allocation of these vaccines is an important research topic [37]. Additionally, there exist individual-based arbovirus models [38] that could be readily applied to perform such analyses. We aim to test our approach on these pathogens as well.

Secondly, we aim to make different algorithmic extensions. In this study, we used elemental bandit learning algorithms (i.e., ϵ -greedy and UCB1). We acknowledge that other algorithms are more suited to identify the optimal action and could potentially learn faster. We created the infrastructure to easily implement and experiment with different algorithms and epidemiological models (details can be found in the Appendices) and we will use this framework

to explore the use of other algorithms. Furthermore, the use of stateless reinforcement learning (i.e., bandits) presents us with a stepping stone to consider reinforcement learning where the partial or full state of the epidemiological model (e.g., which people are currently infected, and which measures have already been taken and to what effect) is used to learn preventive strategies that are more reactive towards events that take place in the simulation. We believe that such strategies may prove to be better than the static strategies we used in this study.

Finally, our current preventive bandits only learn with respect to a single model outcome: more specifically, for influenza this is the proportion of symptomatic infections. In the context of influenza, and for many infectious diseases, there is often interest to consider additional model outcomes (e.g., morbidity, mortality, cost). In the future, we aim to use *multi-objective multi-armed bandits* [39] in contrast to the current single-objective preventive bandits. With this approach, we plan to learn a *coverage set* containing an optimal strategy for every possible preference profile the decision makers might have [40]. We aim to design suitable quality metrics [41–43] tailored to the use case of epidemiological preventive strategy learning, to support the entire spectrum of epidemiological models and thus to prevent method over-fitting [43].

Acknowledgments. Pieter Libin was supported by a PhD grant of the FWO (Fonds Wetenschappelijk Onderzoek – Vlaanderen) and the VUB research council (VUB/OZR2714). Timothy Verstraeten was supported by a PhD grant of the FWO (Fonds Wetenschappelijk Onderzoek – Vlaanderen) and the VUB research council (VUB/OZR2884). Kristof Theys was supported by a postdoctoral grant of the FWO (Fonds Wetenschappelijk Onderzoek – Vlaanderen). Diederik Roijers was supported by a postdoctoral grant of the FWO (Fonds Wetenschappelijk Onderzoek – Vlaanderen). Thanks to Roxana Rădulescu, for her careful proofreading and helpful suggestions. Thanks to Karen Goedeweck, for her careful proofreading. Thanks to all GNU/Python/R/Scala/FluTE hackers for their efforts, useful libraries and excellent tools. The computational resources and services used in this work were provided by the Hercules Foundation and the Flemish Government department EWI-FWO Krediet aan Navorsers (Theys, KAN2012 1.5.249.12.).

Appendix

A FluTE Source

FluTE is a stochastic individual-based model, that is implemented in C++. The original source code, as release by FluTE’s author (i.e., D. Chao), is available from <https://github.com/dlchao/FluTE>. This github repository contains FluTE’s C++ source code, GNU/Linux-specific make files and a set of population density descriptions that can be used to simulate particular geographical settings (i.e., 2000-individual population, Seattle, Los Angeles and the entire United States).

Some changes were made to the source code to make our research easier: we organized the source code in a directory structure and added a CMake meta-make file. This CMake build file allows us to build the source code on GNU/Linux

and MacOS. These changes are publicly available on the <https://github.com/vub-ai-lab/FluTE-bandits> github repository.

B FluTE Configurations

To run our experiments, we defined a model environment to evaluate pre-vaccination with little vaccine available, as described in detail in Sect. 4. The pre-vaccination configuration script can be found in the ‘configs/bandits’ directory of the <https://github.com/vub-ai-lab/FluTE-bandits> github repository. Note that this configuration script is a python Mako template (<http://makotemplates.org/>), to enable easy parameterization of the configuration script.

C Bandit Implementation

We implemented a flexible bandit framework in Scala, the code is publicly available on github: <https://github.com/vub-ai-lab/scala-bandits>. This framework is specifically designed to enable us to easily experiment with new algorithms and environments (i.e., both Scala environments and external environments, such as e.g., the FluTE simulator environment). The repository contains the ϵ -greedy algorithm, the UCB1 algorithm, the Sutton test environment [14], the FluTE environment and some post processing utilities.

D High Performance Computing

Simulating epidemics using individual-based models is a computationally intensive process. Therefore, our experiments were run on a powerful high performance computing cluster: the Flemish Supercomputer Center. We report that, to make this possible, all software had to be installed (or built) for the high performance computing cluster. We report that our FluTE CMake file allows the generation of efficient code (i.e., using SSE instructions) for all platforms used in our analyses (i.e., MacOS, XUbuntu desktop GNU/Linux and GNU/Linux on the high performance computing cluster).

References

1. Stöhr, K.: Influenza: WHO cares. *Lancet Infect. Dis.* **2**(9), 517 (2002)
2. Molinari, N.A.M., Ortega-Sanchez, I.R., Messonnier, M.L., Thompson, W.W., Wortley, P.M., Weintraub, E., Bridges, C.B.: The annual impact of seasonal influenza in the US: measuring disease burden and costs. *Vaccine* **25**(27), 5086–5096 (2007)
3. Nicholls, H.: Pandemic influenza: the inside story. *PLoS Biol.* **4**(2), e50 (2006)
4. Patterson, K.D., Pyle, G.F.: The geography and mortality of the 1918 influenza pandemic. *Bull. Hist. Med.* **65**(1), 4 (1991)

5. Basta, N.E., Chao, D.L., Halloran, M.E., Matrajt, L., Longini, I.M.: Strategies for pandemic and seasonal influenza vaccination of schoolchildren in the United States. *Am. J. Epidemiol.* **170**(6), 679–686 (2009)
6. Germann, T.C., Kadau, K., Longini, I.M., Macken, C.A.: Mitigation strategies for pandemic influenza in the United States. *Proc. Nat. Acad. Sci.* **103**(15), 5935–5940 (2006)
7. Yang, Y., Sugimoto, J.D., Halloran, M.E., Basta, N.E., Chao, D.L., Matrajt, L., Potter, G., Kenah, E., Longini, I.M.: The transmissibility and control of pandemic influenza A (H1N1) virus. *Science* **326**(2009), 729–33 (2009). New York
8. Ajelli, M., Merler, S., Fumanelli, L., y Piontti, A.P., E Dean, N., Longini, I.M., Halloran, M.E., Vespignani, A.: Spatiotemporal dynamics of the Ebola epidemic in Guinea and implications for vaccination and disease elimination: a computational modeling analysis. *BMC Med.* **14**(1), 1–10 (2016)
9. Kraemer, M.U.G., Faria, N.R., Reiner, R.C., Golding, N., Nikolay, B., Stasse, S., Johansson, M.A., Salje, H., Faye, O., William Wint, G.R., et al.: Spread of yellow fever virus outbreak in Angola and the Democratic Republic of the Congo 2015–16: a modelling study. *Lancet Infect. Dis.* **17**(3), 330–338 (2017)
10. Chao, D.L., Halloran, M.E., Obenchain, V.J., Longini Jr., I.M.: FluTE, a publicly available stochastic influenza epidemic simulation model. *PLoS Comput. Biol.* **6**(1), e1000656 (2010)
11. Eubank, S.G., Kumar, V.S., Marathe, M.V., Srinivasan, A., Wang, N.: Structure of social contact networks and their impact on epidemics. *DIMACS Seri. Discrete Math. Theor. Comput. Sci.* **70**(0208005), 181 (2006)
12. Meyers, L.A., Newman, M.E.J., Martin, M., Schrag, S.: Applying network theory to epidemics: control measures for *Mycoplasma pneumoniae* outbreaks. *Emerg. Infect. Dis.* **9**(2), 204–210 (2003)
13. Fumanelli, L., Ajelli, M., Merler, S., Ferguson, N.M., Cauchemez, S.: Model-based comprehensive analysis of school closure policies for Mitigating Influenza Epidemics and Pandemics. *PLoS Comput. Biol.* **12**(1), e1004681 (2016)
14. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge (1998)
15. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* **47**(2–3), 235–256 (2002)
16. World Health Organization, et al.: WHO guidelines on the use of vaccines and antivirals during influenza pandemics (2004)
17. Enserink, M.: Crisis underscores fragility of vaccine production system. *Science* **306**(5695), 385 (2004)
18. Biggerstaff, M., Reed, C., Swerdlow, D.L., Gambhir, M., Graitcer, S., Finelli, L., Borse, R.H., Rasmussen, S.A., Meltzer, M.I., Bridges, C.B.: Estimating the potential effects of a vaccine program against an emerging influenza pandemic - United States. *Clin. Infect. Dis.* **60**, S20–S29 (2015)
19. Miller, M.A., Viboud, C., Balinska, M., Simonsen, L.: The signature features of influenza pandemics: implications for policy. *N. Engl. J. Med.* **360**(25), 2595–2598 (2009)
20. Garnett, G.P., Cousens, S., Hallett, T.B., Steketee, R., Walker, N.: Mathematical models in the evaluation of health programmes. *Lancet* **378**(9790), 515–525 (2011)
21. Lessler, J., Edmunds, W.J., Halloran, M.E., Hollingsworth, T.D., Lloyd, A.L.: Seven challenges for model-driven data collection in experimental and observational studies. *Epidemics* **10**, 78–82 (2014)

22. Rahmandad, H., Sterman, J.: Heterogeneity and network structure in the dynamics of diffusion: comparing agent-based and differential equation models. *Manage. Sci.* **54**(5), 998–1014 (2008)
23. Willem, L.: Agent-Based Models For Infectious Disease Transmission: Exploration, Estimation & Computational Efficiency. Ph.D. thesis (2015)
24. Halloran, M.E., Longini, I.M., Nizam, A., Yang, Y.: Containing bioterrorist smallpox. *Science* **298**(5597), 1428–1432 (2002). New York
25. Herbert, R.: Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.* **58**(5), 527–535 (1952)
26. Press, W.H.: Bandit solutions provide unified ethical models for randomized clinical trials and comparative effectiveness research. *Proc. Nat. Acad. Sci.* **106**(52), 22387–22392 (2009)
27. Gittins, J., Glazebrook, K., Weber, R.: *Multi-Armed Bandit Allocation Indices*, 2nd edn. Wiley, Chichester (2011)
28. Awerbuch, B., Kleinberg, R.: Near-optimal adaptive routing: shortest paths and geometric generalizations. In: *Proceeding of the 36th Annual ACM Symposium on Theory of Computing*, pp. 45–53 (2004)
29. Kaufmann, E., Cappé, O., Garivier, A.: On the complexity of A/B testing. In: *COLT*, pp. 461–481 (2014)
30. Blum, A., Kumar, V., Rudra, A., Felix, W.: Online learning in online auctions. *Theoret. Comput. Sci.* **324**(2–3), 137–146 (2004)
31. Medlock, J., Galvani, A.P.: Optimizing influenza vaccine distribution. *Science* **325**(5948), 1705–1708 (2009)
32. McLean, H.Q., Thompson, M.G., Sundaram, M.E., Kieke, B.A., Gaglani, M., Murthy, K., Piedra, P.A., Zimmerman, R.K., Nowalk, M.P., Raviotta, J.M., Jackson, M.L., Jackson, L., Ohmit, S.E., Petrie, J.G., Monto, A.S., Meece, J.K., Thaker, S.N., Clippard, J.R., Spencer, S.M., Fry, A.M., Belongia, E.A.: Influenza vaccine effectiveness in the United States during 2012–2013: variable protection by age and virus type. *J. Infect. Dis.* **211**(10), 1529–1540 (2015)
33. Abbas, A.K., Lichtman, A.H., Pillai, S.: *Cellular and Molecular Immunology*. Elsevier Health Sciences, Amsterdam (2014)
34. CDC: Key facts about influenza (flu) & flu vaccine. Centers for Disease Control and Prevention, Atlanta (2014)
35. Hadinegoro, S.R., Arredondo-García, J.L., Capeding, M.R., Deseda, C., Chotpitayasunondh, T., Dietze, R., Muhammad Ismail, H.I.H., Reynales, H., Limkittikul, K., Rivera-Medina, D.M., Tran, H.N., Bouckennooghe, A., Chansinghakul, D., Cortés, M., Fanouillere, K., Forrat, R., Frago, C., Gailhardou, S., Jackson, N., Noriega, F., Plennevaux, E., Wartel, T.A., Zambrano, B., Saville, M.: Efficacy and long-term safety of a dengue vaccine in regions of endemic disease. *N. Engl. J. Med.* **373**(13), 1195–1206 (2015)
36. Cohen, J.: The race for a Zika vaccine is on. *Science* **351**(6273), 543–544 (2016)
37. Ferguson, N.M., Rodríguez-Barraquer, I., Dorigatti, I., Mier-y Teran-Romero, L., Laydon, D.J., Cummings, D.A.T.: Benefits and risks of the Sanofi-Pasteur dengue vaccine: modeling optimal deployment. *Science* **353**(6303), 1033–1036 (2016)
38. Chao, D.L., Halstead, S.B., Halloran, M.E., Longini, I.M.: Controlling dengue with vaccines in Thailand. *PLoS Neglected Trop. Dis.* **6**(10), e1876 (2012)
39. Drugan, M.M., Nowe, A.: Designing multi-objective multi-armed bandits algorithms: a study. In: *Proceedings of the International Joint Conference on Neural Networks* (2013)
40. Roijers, D.M., Vamplew, P., Whiteson, S., Dazeley, R.: A survey of multi-objective sequential decision-making. *J. Artif. Intell. Res.* **48**, 67–113 (2013)

41. Roijers, D.M.: Multi-objective decision-theoretic planning. Ph.D. thesis, University of Amsterdam (2016)
42. Vamplew, P., Dazeley, R., Berry, A., Issabekov, R., Dekker, E.: Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Mach. Learn.* **84**(1–2), 51–80 (2011)
43. Zintgraf, L.M., Kanters, T.V., Roijers, D.M., Oliehoek, F.A., Beau, P.: Quality assessment of MORL algorithms: a utility-based approach. In: *Proceedings of the Twenty-Fourth Belgian-Dutch Conference on Machine Learning, Benelearn 2015* (2015)