

Finding optimal vaccination strategies for pandemic influenza using genetic algorithms

Rajan Patel*, Ira M. Longini Jr., M. Elizabeth Halloran

Department of Biostatistics, The Rollins School of Public Health, Emory University, 1518 Clifton Road, Atlanta, GA 30322, USA

Received 6 May 2004; received in revised form 23 September 2004; accepted 22 November 2004

Available online 20 January 2005

Abstract

In the event of pandemic influenza, only limited supplies of vaccine may be available. We use stochastic epidemic simulations, genetic algorithms (GA), and random mutation hill climbing (RMHC) to find optimal vaccine distributions to minimize the number of illnesses or deaths in the population, given limited quantities of vaccine. Due to the non-linearity, complexity and stochasticity of the epidemic process, it is not possible to solve for optimal vaccine distributions mathematically. However, we use GA and RMHC to find near optimal vaccine distributions. We model an influenza pandemic that has age-specific illness attack rates similar to the Asian pandemic in 1957–1958 caused by influenza A(H2N2), as well as a distribution similar to the Hong Kong pandemic in 1968–1969 caused by influenza A(H3N2). We find the optimal vaccine distributions given that the number of doses is limited over the range of 10–90% of the population. While GA and RMHC work well in finding optimal vaccine distributions, GA is significantly more efficient than RMHC. We show that the optimal vaccine distribution found by GA and RMHC is up to 84% more effective than random mass vaccination in the mid range of vaccine availability. GA is generalizable to the optimization of stochastic model parameters for other infectious diseases and population structures.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Genetic algorithms; Influenza; Stochastic models; Optimization

1. Introduction

Influenza is a major public health concern. Influenza spreads rapidly in seasonal epidemics which cost society a considerable amount in terms of health care expenses, lost productivity, and loss of life. Globally, influenza annually costs from \$71 and \$167 billion and results in 250 000 and 500 000 deaths (World Health Organization, 2004). Annual influenza epidemics occur partially due to strains of influenza genetically drifting from year to year. The influenza vaccine produced before the influenza season is targeted against the strains that are predicted to circulate in the coming season. However, a major antigenic shift can occur with little warning, resulting in pandemic

influenza (Kilbourne, 1975). The two most recent pandemics were the “Asian Flu” A (H2N2) of 1957–1958 (Elveback et al., 1976; Jordan, 1961; Longini et al., 1978) and “Hong Kong Flu” A (H3N2) of 1968–1969 (Davis et al., 1970; Elveback et al., 1976; Longini et al., 1978; Sharrar, 1969). The Asian Flu is estimated to have caused 70 000 deaths in the US, while the Hong Kong Flu is estimated to have caused 34 000 deaths in the US. In addition, the Hong Kong pandemic of 1968–1969 is estimated to have cost society 3.8 billion dollars in the US alone (Kavet, 1972). Should a major antigenic shift occur, there may be time to produce only a limited amount of vaccine efficacious against the new strain and ensuing pandemic. Knowledge on how to distribute the limited supply of vaccine optimally among different age groups would help us to minimize the impact of the epidemic. This impact can be measured in many ways, two of which are number of illnesses and loss of life.

*Corresponding author. Tel.: +1 404 727 9169;
fax: +1 404 727 1370.

E-mail address: rspate2@emory.edu (R. Patel).

Optimization methods have been developed for deterministic simulation models (Anderson and May, 1991; Greenhalgh, 1986; Hethcote and Waltman, 1973; Wickwire and Guest, 1976), and optimization studies have been done to find the vaccine distributions for influenza using a simple deterministic model (Longini et al., 1978). However, they have not been done with more complex stochastic simulation models. Influenza is transmitted in a complex way from person to person. In addition, given an introduction of influenza into a population, the probability of a major epidemic and the possible size of an epidemic are highly variable. Thus, the mathematical models for influenza epidemics should have a detailed contact structure and be stochastic. In addition, the epidemic process is non-linear since the incidence of new infections depends on the current number of both infectives and susceptibles in the population at a particular time. All of these factors make optimization based on traditional gradient methods, such as the Newton–Raphson method, difficult or even prohibitive. Robbins and Munro (1951) developed a stochastic approximation method whose convergence is guaranteed under mild conditions. The method, however, requires knowledge of the analytic gradient of the considered objective function (Weisstein, 2004). Kiefer and Wolfowitz (1952) developed an extension to the Robbins–Munro algorithm. However, in terms of simulation optimization, the drawback to both of these methods remains the unavailability of gradients. In the case of our stochastic simulation multi-dimensional optimization problem, we consider genetic algorithms (GA) and random mutation hill climbing (RMHC) as stochastic optimization methods.

In this paper, we find optimal distributions of a limited supply of vaccine in the event of pandemic influenza generated by a stochastic simulation model using GA (Holland, 1975) as well as RMHC (Forrest and Mitchell, 1993). We configure the model to simulate the baseline illness attack rates consistent with the past patterns of Asian and Hong Kong pandemic influenza. We find optimal vaccine distributions in terms of minimizing influenza illness or death. These optimization methods are generally applicable to other infectious diseases and population structures.

2. Methods

2.1. The simulation model

We use a discrete time, stochastic simulation model of influenza spread within a structured population of 10000 individuals to estimate the effectiveness of various distributions of a given amount of vaccine. This model is a direct extension of an earlier model developed for influenza intervention studies (Halloran et al., 2002a;

Longini et al., 2004). The model simulates the stochastic spread of influenza in a population where the age and family structure approximate information from the US Census 2000. The 10000 person population consists of five 2000 person communities each containing four neighborhoods, one high school, one middle school, and two elementary schools (Fig. 1A). Pre-school children visit either a small play group or a large day care center. Working adults make contact in workplaces. Individuals can thus come into contact with other members of their own family (household contacts), people at school or work, and others in their neighborhood and community,




Fig. 1. Structure of the populations. (A) The 2000 person communities consist of households embedded in neighborhoods. Each community is partitioned into four neighborhoods. Small children mix in playgroups and daycare centers within their neighborhoods. The school mixing groups link neighborhoods as shown. (B) Five communities are linked to create a larger population of size 10000 people. This linkage is done by allowing adults who work to be randomly assigned to work in mixing groups of size 25 throughout the whole population.

each group with its own transmission probability depending upon the age of the individual. Influenza can spread among communities by workplace contacts (Fig. 1B). Each day, for each susceptible individual, the probability of becoming infected is calculated based on that individual’s vaccination status, the vaccination status of the individual’s contacts, and the group-specific transmission probabilities. Influenza is introduced into the population by randomly assigning 12 initial infectives in each of the five communities. The initial infectives are omitted from attack rate calculations.

People who become infected enter the latent period and incubation period (1.9 days mean length) during which they are not infectious and do not show symptoms. This is followed by the infectious period (4.1 day mean length) during which they may show influenza symptoms. In the model, if an individual shows symptoms, he may withdraw to the home after 1 day of infectiousness with some probability, then coming into contact with only other members of his household. More details on the model can be found in previous publications (Halloran et al., 2002b; Longini et al., 2004).

We run the optimization algorithm with two basic attack rate patterns. We calibrate the baseline epidemic (all individuals are unvaccinated) to have age-specific illness attack rates similar to the 1957–1958 A (H2N2) Asian influenza pandemic. We will refer to this as Asian-like influenza for the rest of this paper. We also calibrate the baseline epidemic to the 1968–1969 A (H3N2) Hong Kong influenza pandemic. We will refer to this as Hong Kong-like influenza for the rest of this paper. Table 1 shows the illness attack rates from the literature. Asian-like influenza had the highest illness attack rate in school children, followed by preschool children with adults

having a lower attack rate. In contrast, Hong Kong-like influenza had a relatively flat attack rate across age groups.

In our simulations, vaccination of the population occurs before the influenza pandemic begins so that vaccinated people will develop immunity. Thus, all of the available doses of the vaccine will be administered before influenza begins to spread in the population. We assume that the vaccine efficacy for susceptibility is $VE_S = 0.70$ and the vaccine efficacy for infectiousness is $VE_I = 0.80$ (Belshe et al., 1998, 2000a, b; Longini et al., 2000).

2.2. The optimization problem

The optimization problem is as follows: Given a limited quantity of influenza vaccine and a particular population structure and illness attack rate pattern for a single wave of pandemic influenza, what proportion of each age group should be vaccinated to minimize the impact of the epidemic?

We divide the population into five age groups: preschool, school, young adults, middle aged adults, and old adults, that are indexed as $i = 1, \dots, 5$. We let n_i be the number of individuals in age group i , and the total population size is $n = \sum_{i=1}^5 n_i$. We let σ_i be the average illness attack rate in age group i , i.e. proportion of both vaccinated and unvaccinated people in age group i that get infected and become ill over the course of the epidemic. The average is over a predetermined number of simulated series for each vaccination distribution. We let V be the total number of vaccine doses available before day one of the epidemic, and v_i the proportion of age group i that is vaccinated before day one of the epidemic. Then, $Q = \sum_{i=1}^5 n_i v_i$ is the total number of doses distributed, where $Q \leq V$, since we cannot use more vaccine than there is available. We assume that each person vaccinated receives one dose of vaccine. Also, define the array $v = (v_1, \dots, v_5)$ as the control vector. We refer to v_i/n_i as the vaccine coverage in age group i , and V/n as the total vaccine coverage for the population. To reflect the impact of a single illness, we let w_i be the weight assigned to an illness in each age group for minimization of the loss function. Then, the optimization problem is expressed as the minimization of

$$\sum_{i=1}^5 \sigma_i n_i w_i, \tag{1}$$

such that

$$\sum_{i=1}^5 v_i n_i \leq V. \tag{2}$$

We concentrate on minimizing overall illness in the population as well as number of lives lost given a predetermined number of doses V of vaccine. We find

Table 1
Model specifications

Age groups	Pre-school	School	Young adults	Middle adults	Old adults
Age (years)	0–4	5–18	19–50	51–64	65+
Group size (%) (10 000 people)	6.80	20.40	46.24	14.09	12.47
Deaths per 10 000 illnesses ^a	0.263	0.210	2.942	2.942	199.8
Target illness attack rates (%)					
Asian-like influenza ^b	35.2	55.4	24.6	19.9	13.9
Hong Kong- like influenza ^c	34.8	34.8	34.6	32.2	30.5

^aLongini et al., 2004; Thompson et al., 2003.

^bElveback et al., 1976; Jordan, 1961; Longini et al., 1978.

^cDavis et al., 1970; Elveback et al., 1976; Longini et al., 1978; Sharrar, 1969.

unique optimal distributions of the vaccine for varying numbers of doses of available vaccine, V as well as different weights, w_i , based on our objective function. Thus, for each combination of V and set of weights, w_i , we obtain and solve a unique optimization problem. We will refer to the control vector v that minimizes (1) as the optimal vaccine distribution for the given number of doses V and weights w_i .

2.3. The genetic algorithm

A GA is an algorithm which can find approximate solutions by using an approach inspired by the biological processes of reproduction natural selection. The general GA will begin with an initial set of randomly generated individuals. An individual is a set of parameters that are a candidate solution to the optimization problem and can be evaluated by some fitness function. We can refer to the parameters as genes. The individuals in each generation are evaluated and sorted based on a fitness function. A second pool of individuals is generated by repeatedly selecting a pair of good individuals from the previous generation and breeding them using the biological principles of inheritance and mutation. This subsequent generation of individuals can then be evaluated by the fitness function, and the process can repeat. There are many variations to the typical GA, and we introduce a simple implementation which will render near-optimal vaccine strategies given a stochastic simulation model.

We define an individual, v_G , which consists of the parameters,

$$\{v_{G1}, v_{G2}, v_{G3}, v_{G4}, v_{G5}\}, \tag{3}$$

as a distribution of the V doses of vaccine among the 5 age groups such that (2) is satisfied at equality. The fitness function we implement calculates the estimated number of illnesses or number of lives lost, depending on our objective, when pre-vaccinating the population according to the vaccine distribution given by v_G .

The algorithm is as follows:

Initialization

1. Randomly generate an initial set of 50 individuals. We generate the initial pool of individuals such that (3) is randomly determined for each individual under the condition that (2) is satisfied at equality.

Iteration

2. Evaluate each of the 50 individuals in the pool by pre-vaccinating the proportion of the population according to the individual being evaluated, then

running the stochastic simulator 20 times. The resulting fitness of each individual is the mean of

$$\sum_{i=1}^5 \sigma_i n_i w_i \tag{4}$$

over the 20 runs of the simulator. A smaller value of this fitness function represents a more fit individual, as this fitness function is essentially a loss function.

3. Sort the 50 individuals according to their fitness.
4. We select the best 25 individuals from the current generation and pass them directly to the next generation of individuals. We do this to ensure that no subsequent generation takes a step backwards and consists of less fit individuals than the previous generation. We generate the remaining 25 individuals by breeding randomly selected pairs from the previous generation. We use a deterministic tournament selection method repeatedly to select maternal and paternal individuals to breed. The breeding process is implemented as follows.
 - (a) We randomly select a paternal individual using the following deterministic tournament selection method with tournament size equal to 10. We randomly select 10 individuals from the current generation and sort them based on their measured fitness. We select the best individual out of those 10 to serve as a paternal individual for purpose of breeding.
 - (b) We then randomly select a maternal individual in the same manner. We disallow the previously selected paternal individual from inclusion in the tournament so that the same individual cannot breed with himself. We are able to increase and decrease selection pressure (bias of selecting better individuals) by increasing and decreasing the tournament size, respectively. After some trials, we determined that a tournament size of 10 generates an adequate amount of selection pressure.
 - (c) We generate a child, v_C , by breeding the two randomly selected individuals above. We denote these parents as v_M and v_F . We use a uniform crossover breeding method with a mixing ratio of 80%. Thus, there is an 80% probability that the child will inherit a specific gene from the father, and 20% probability that the child will inherit that gene from the mother. Inheritance of each gene is independent from any other gene. The 80:20 mixing ratio was suggested as optimal by Spears and De Jong (1990) when using the uniform crossover mixing method.
 - (d) We begin the breeding by assigning:

$$\forall i \quad v_{Ci}^{(1)} = \begin{cases} v_{Mi}, & \text{with probability 0.8,} \\ v_{Fi}, & \text{with probability 0.2.} \end{cases} \tag{5}$$

- (e) Since $v_C^{(1)}$ likely does not satisfy (2) at equality, we must add or subtract doses from randomly selected age groups. This process serves two processes. First it ensures that the child satisfies (2), and secondly it introduces a mutation step into the breeding procedure. Mutation allows genetic diversity to pass from one generation to the next.
- (f) We iteratively update $v_C^{(j)}$ by randomly selecting one age group i at each iteration and updating it in the following manner until $v_C^{(j)}$ satisfies (2) at equality:

$$v_{Ci}^{(j+1)} = \begin{cases} v_{Ci}^{(j)} + \frac{\min(V - \sum_{i=1}^5 v_{Ci}^{(j)} n_i, (1 - v_{Ci}^{(j)}) n_i (V - Q)/5)}{n_i} & \text{if } Q < V, \\ v_{Ci}^{(j)} - \frac{\min(\sum_{i=1}^5 v_{Ci}^{(j)} n_i - V, v_{Ci}^{(j)} n_i (Q - V)/5)}{n_i} & \text{if } Q > V. \end{cases} \quad (6)$$

The inclusion of $(Q - V)/5$ (mutation parameter) into the *min* function of (6) allows us to control how far each of the v_{Ci} can stray from the respective parameter of the parent from which it was obtained. The inclusion of this specific mutation parameter requires about five relatively small repairs to be made to the genes of the new child. The number of repairs is directly related to the denominator of the parameter, while the size of the mutation of each gene is inversely related to the denominator of the mutation parameter. The GA literature provides several methods to perform a constrained crossover breeding including the “repair” algorithm used here. The main drawback of repair algorithms is cost in terms of running time (Orvosh and Davis, 1993). However, the bulk of the running time of our optimization is involved with evaluating each individual, as our influenza simulator takes approximately 2 seconds to run and we run the simulator 20 times in evaluating each individual. Thus the relative time cost of breeding is negligible, thus we elected to use the simpler repair algorithm as described here.

5. We generate 25 children in a similar manner to fill the next generation of individuals. We allow previously selected maternal and paternal individuals to be used again to breed more children.

Convergence

6. We repeat this process from Step 2 to Step 5 until the best individual does not change for four consecutive generations. The best individual in the final generation will yield the optimal vaccine distribution given

the quantity of available influenza vaccine, the particular population structure, the illness attack rate pattern, and the objective of the optimization.

2.4. Random mutation hill climbing

We compare the GA approach to another optimization approach, RMHC, which may converge more quickly (Forrest and Mitchell, 1993; Mitchell et al., 1994). The RMHC algorithm finds an optimal individual from an initial randomly generated individual by attempting different modifications of the initial individual and calculating the fitness of each modification in a similar manner to the GA. Once a modification of the individual is found that has a better fitness than the original, that individual replaces the original and modifications of the new individual are considered. We continue until several modifications of the current individual fail to yield a better individual than the current one. The algorithm works in the following manner:

Initialization

1. Generate a random individual v_A as defined in the GA approach which satisfies (2) at equality.
2. Evaluate v_A by pre-vaccinating the proportion of the population according to v_A running the stochastic simulator 20 times. The resulting fitness of v_A is the mean of

$$\sum_{i=1}^5 \sigma_i n_i w_i \quad (7)$$

over the 20 runs of the simulator.

Iteration

3. Randomly mutate v_A by first creating a copy of v_A , $v_B^{(1)}$, and mutating $v_B^{(1)}$ in the following manner:
 - (a) Select one age group, i , at random, and set

$$v_{Bi}^{(1)} = v_{Ai} + U, \quad (8)$$

where U is randomly chosen from a Uniform(0, $1 - v_{Ai}$) distribution. A Gaussian distribution is commonly used to mutate continuous genes, however we noticed no difference in the quality of resulting individuals and running time of the algorithm using the uniform distribution. The advantage of using a uniform distribution in this case is that its possible values are bounded so that we do not require artificial clipping to ensure a gene remains in the $[0, 1]$ range. We must now “repair” the individual to satisfy (2). The repair method is unnecessary in

Forrest and Mitchell's RMHC algorithm (Forrest and Mitchell, 1993) as they dealt with unconstrained optimization.

(b) We iteratively update $\mathbf{v}_B^{(j)}$ by randomly selecting one age group i at each iteration and updating it in the following manner until $\mathbf{v}_B^{(j)}$ satisfies (2) at equality. Let $Q = \sum_{i=1}^5 v_{Bi}^{(1)} n_i$.

$$v_{Bi}^{(j+1)} = v_{Bi}^{(j)} - \frac{\min(\sum_{i=1}^5 v_{Bi}^{(j)} n_i - V, v_{Bi}^{(j)} n_i, (Q - V)/5)}{n_i}. \quad (9)$$

(c) The result of this process yields \mathbf{v}_B which satisfies (2) at equality and thus is a valid individual.

4. Determine the fitness of \mathbf{v}_B by calculating the mean of (4) over 20 runs of the simulator under \mathbf{v}_B .
5. If \mathbf{v}_B is as fit as or more fit than \mathbf{v}_A , then we replace \mathbf{v}_A with \mathbf{v}_B and begin the mutation process again. If \mathbf{v}_A is more fit than \mathbf{v}_B , we retain \mathbf{v}_A and try another mutation.

Convergence

6. We continue this process until 100 consecutive mutations of \mathbf{v}_A do not reveal a more fit individual than \mathbf{v}_A .

2.5. Parameter values

We calibrated the simulation model for two different potential pandemic influenza illness attack rate patterns due to Asian-like and Hong Kong-like influenza. We varied the transmission probabilities to produce roughly the illness attack rates shown in the Table 1.

We ran the optimization routines for a range of V such that we had 10–90% overall vaccine coverage in the population. We ran the optimization minimizing overall illness as well as minimizing overall death. We use weights of one, i.e. $w_i = 1, i = 1, \dots, 5$, for the minimizing illness objective function, and we used the influenza death rates, i.e. $w_1 = 0.263, w_2 = 0.210, w_3 = 2.942, w_4 = 2.942, w_5 = 199.8$, for the minimizing death objective function. When minimizing illness, the illnesses in each age group have the same weight. However when minimizing death, we weight each illness by the approximate number of deaths per 10 000 illnesses in each age group (Longini et al., 2004). This varies significantly such that the death rate in older adults is almost 1000 times as large as that of children.

3. Results

3.1. Optimal vaccine distributions

Tables 2–5 show the results given by the GA. The results from RMHC are similar to those from the GA,

and are thus not in the tables. For example, the RMHC result for the Asian-like flu case when minimizing illness attack rate and $V/n = 0.40$ is $\mathbf{v} = (0.88, 1.00, 0.21, 0.26, 0.00)$. We do not expect exact results from each algorithm as the results are approximations of the optimal distributions. The two algorithms converging to significantly different distributions of the vaccine may suggest different local optima, however, we did not encounter this case.

Table 2 gives optimal vaccine distributions for an Asian-like influenza epidemic when minimizing the illness attack rate. When no vaccine is available ($V/n = 0$), the simulated illness attack rate pattern is similar to the target baseline given in Table 1. When we have enough vaccine to vaccinate 10% of the population ($V/n = 0.10$), the GA and RMHC optimization algorithms converge to an optimal vaccine distribution where only school children should be vaccinated. With 20% vaccine coverage, we can effectively stop the epidemic as we are able to vaccinate 98% of school children, surpassing the critical vaccination fraction as described in Hill and Longini (2003). With coverages higher than 20%, it is optimal to begin vaccinating pre-school children, followed by young and middle aged adults, and finally older adults. After a vaccine coverage of 20%, we experience diminished returns for each additional dose of vaccine available. Although the spread of the influenza agent has effectively stopped, these additional doses of vaccine still prevent some additional cases of influenza.

Table 3 gives optimal vaccine distributions when minimizing the overall number of deaths in the population. When there is enough vaccine to vaccinate 10% of the population, the optimization algorithms suggest vaccinating only older adults. However, when $V/n = 0.20$, we are able to achieve a vaccine distribution which can prevent the epidemic. This distribution suggests vaccinating mainly school-aged children, as they are the most potent group of Asian-like influenza spreaders. As the amount of available vaccine increases, we begin to vaccinate the older adults, followed by pre-school children and finally young adults and middle-aged adults.

Table 4 gives optimal vaccine distributions as found by the GA under the scenario of a Hong Kong-like influenza pandemic and with the goal of minimizing the illness attack rate. At a vaccine coverage of 10% ($V/n = 0.10$), the optimization algorithms suggest vaccination of school-age children and young adults. As more vaccine becomes available, pre-school children and middle-age adults should be vaccinated, and finally, older adults should receive vaccine. Unlike the optimal vaccination distributions in the Asian-like influenza case, the optimal vaccine distributions in the Hong Kong-like influenza case are more spread out among the age groups, as no single group is responsible for a

Table 2
Optimal vaccine distributions for minimizing illness (Asian-like influenza)

Coverage (%)	Optimal vaccine distribution					Illness attack rate					
	PS ^a	S	YA	MA	OA	PS	S	YA	MA	OA	AVE ^b
0	0.00	0.00	0.00	0.00	0.00	0.338	0.559	0.238	0.223	0.160	0.299
10	0.00	0.49	0.00	0.00	0.00	0.182	0.221	0.127	0.118	0.081	0.143
20	0.00	0.98	0.00	0.00	0.00	0.144	0.134	0.101	0.092	0.065	0.105
30	0.51	1.00	0.10	0.12	0.00	0.021	0.012	0.025	0.022	0.017	0.021
40	0.92	1.00	0.22	0.22	0.00	0.008	0.010	0.016	0.015	0.012	0.014
50	0.94	1.00	0.35	0.46	0.07	0.007	0.009	0.011	0.009	0.009	0.010
60	1.00	1.00	0.48	0.69	0.09	0.005	0.009	0.008	0.006	0.008	0.008
70	1.00	1.00	0.73	0.61	0.03	0.004	0.007	0.005	0.005	0.006	0.006
80	1.00	1.00	0.86	0.80	0.13	0.004	0.007	0.003	0.004	0.005	0.005
90	1.00	1.00	1.00	0.97	0.25	0.004	0.007	0.002	0.002	0.004	0.004
100	1.00	1.00	1.00	1.00	1.00	0.004	0.007	0.002	0.002	0.002	0.003

^aPS, S, YA, MA, OA abbreviate the 5 age-groups defined in Table 1.

^bAVE represents the average rate over 100 simulations given the specified vaccine distribution.

Table 3
Optimal vaccine distributions for minimizing death (Asian-like influenza)

Coverage (%)	Optimal vaccine distribution					Death rate (×10000)					
	PS	S	YA	MA	OA	PS	S	YA	MA	OA	AVE
0	0.00	0.00	0.00	0.00	0.00	0.089	0.117	0.701	0.657	32.02	4.440
10	0.00	0.00	0.00	0.00	0.80	0.085	0.116	0.652	0.605	13.57	2.108
20	0.00	0.86	0.00	0.03	0.15	0.017	0.006	0.133	0.120	5.091	0.716
30	0.00	0.88	0.01	0.22	0.68	0.013	0.005	0.099	0.078	2.216	0.335
40	0.32	1.00	0.06	0.17	1.00	0.007	0.002	0.066	0.056	0.837	0.144
50	0.48	0.93	0.19	0.45	1.00	0.004	0.003	0.047	0.035	0.667	0.111
60	1.00	1.00	0.23	0.69	1.00	0.001	0.002	0.034	0.019	0.521	0.084
70	0.92	1.00	0.39	0.92	1.00	0.002	0.002	0.025	0.012	0.450	0.070
80	1.00	1.00	0.62	0.84	1.00	0.001	0.002	0.016	0.011	0.376	0.056
90	1.00	1.00	0.85	0.78	1.00	0.001	0.002	0.010	0.010	0.328	0.047
100	1.00	1.00	1.00	1.00	1.00	0.001	0.001	0.007	0.006	0.318	0.044

Table 4
Optimal vaccine distributions for minimizing illness (Hong Kong-like influenza)

Coverage (%)	Optimal vaccine distribution					Illness attack rate					
	PS	S	YA	MA	OA	PS	S	YA	MA	OA	AVE
0	0.00	0.00	0.00	0.00	0.00	0.343	0.352	0.352	0.337	0.278	0.340
10	0.00	0.19	0.13	0.00	0.00	0.263	0.226	0.247	0.253	0.212	0.240
20	0.12	0.40	0.20	0.13	0.00	0.154	0.116	0.154	0.154	0.135	0.144
30	0.00	0.28	0.44	0.28	0.00	0.098	0.078	0.063	0.068	0.068	0.070
40	0.28	0.53	0.49	0.34	0.00	0.039	0.028	0.032	0.034	0.035	0.032
50	0.33	0.63	0.58	0.57	0.00	0.025	0.015	0.017	0.017	0.022	0.018
60	0.65	0.79	0.74	0.37	0.00	0.012	0.008	0.010	0.015	0.016	0.011
70	0.83	0.89	0.76	0.80	0.00	0.007	0.006	0.008	0.006	0.013	0.008
80	0.88	0.90	0.96	0.80	0.00	0.006	0.005	0.004	0.005	0.011	0.006
90	1.00	1.00	1.00	1.00	0.20	0.004	0.003	0.003	0.003	0.008	0.004
100	1.00	1.00	1.00	1.00	1.00	0.003	0.003	0.003	0.003	0.002	0.003

Table 5
Optimal vaccine distribution for minimizing death (Hong Kong-like influenza)

Coverage (%)	Optimal vaccine distribution					Death rate ($\times 10000$)					
	PS	S	YA	MA	OA	PS	S	YA	MA	OA	AVE
0	0.00	0.00	0.00	0.00	0.00	0.090	0.074	1.039	0.992	55.56	7.570
10	0.00	0.00	0.00	0.00	0.80	0.080	0.066	0.900	0.862	22.66	3.383
20	0.00	0.16	0.08	0.05	1.00	0.060	0.043	0.644	0.622	11.45	1.827
30	0.00	0.35	0.10	0.41	1.00	0.036	0.021	0.374	0.274	6.172	0.988
40	0.13	0.36	0.26	0.53	1.00	0.017	0.011	0.163	0.119	3.015	0.472
50	0.25	0.53	0.42	0.41	1.00	0.009	0.005	0.078	0.072	1.724	0.263
60	0.56	0.53	0.51	0.67	1.00	0.004	0.003	0.047	0.037	1.047	0.158
70	0.09	0.62	0.74	0.70	1.00	0.006	0.003	0.028	0.027	0.897	0.130
80	0.70	0.84	0.73	0.84	1.00	0.002	0.001	0.022	0.016	0.751	0.106
90	1.00	0.61	1.00	0.86	1.00	0.001	0.002	0.011	0.014	0.551	0.076
100	1.00	1.00	1.00	1.00	1.00	0.001	0.001	0.010	0.009	0.448	0.062

majority of the illnesses. We are unable to obtain an optimal vaccination distribution to prevent an epidemic until we have approximately 40% coverage, as opposed to only 20% coverage in the Asian-like influenza case. We begin seeing increasingly diminished returns on each additional dose of vaccine after about 40 or 50% coverage.

Table 5 gives optimal vaccine distributions for a Hong Kong-like influenza when we minimize the number of deaths. The optimal vaccination distributions suggest vaccinating older adults when coverage is low ($V/n = 0.10$). As more vaccine becomes available, it is optimal to give vaccine to school-children, young adults, and middle-age adults relatively evenly. We do not experience a shift in the age-group for which it is optimal to vaccinate as we do in the Asian-like influenza case when minimizing the number of deaths. This difference occurs because in the Hong Kong-like influenza setting, we are unable to stop the epidemic early with high vaccine coverage in one of the age-groups because the virus spreads more homogeneously than it does in the Asian-like influenza setting.

The effectiveness of the optimal vaccine distributions for Asian-like influenza when compared to random mass vaccination of the entire population is given in Table 6. For a given overall vaccine coverage level, we compare the overall illness attack rates and death rates for the optimal vaccination distribution with those if the vaccine were randomly distributed. We use the relative overall effectiveness comparing the two vaccination strategies as our measure of how much better the optimal vaccine distribution is than is random vaccination. Thus, we define the relative overall effectiveness (Halloran et al., 1999) as $= 1 - \frac{index_{opt}}{index_{ran}}$, where $index_{opt}$ is either the illness attack rate or death rate when the optimal vaccination distribution is used and $index_{ran}$ is the complementary index when random vaccination is used. For example, from Table 6, implementation of the

optimal vaccine distribution given 30% coverage, is 84% more effective than random vaccination. For both objective functions, the effectiveness of the optimal vaccine distributions compared to random vaccination is highest at coverage levels 20–70% and peaks around 30–40%.

3.2. Algorithm convergence

Neither GA nor RMHC optimization algorithms are guaranteed to converge to exactly the optimal vaccine distribution. The randomness of each algorithm and the difficulty of the optimization problem allows us to expect only an approximation of the optimal vaccine distribution. The results we present in this analysis may or may not be optimal in the sense that there exists no better distribution of the vaccine. However, we present approximations of the optimal distribution suggested by these algorithms. It is unlikely that vaccination of an age group not suggested by any of the optimizations will result in a better distribution of the vaccine.

Both algorithms did converge to a similar distribution of the vaccine, however GA converged much more quickly than did RMHC for almost all of the possible vaccine coverages, all optimization criteria, and for both types of influenza strain that we considered. Table 7 shows the number of simulations for each algorithm to converge in the Asian-like influenza case for one run of the algorithm for each of the vaccine coverages considered. Each run of the stochastic simulation takes approximately 2 seconds on a Pentium 4 1.6 GhZ machine with 256 MB RAM. We run the simulation 20 times in evaluating the fitness of each individual, each individual takes approximately 40 seconds to evaluate. Evaluation of the individuals was the overwhelmingly predominant time cost with both the GA and RMHC algorithms.

Table 6
Asian-like influenza optimal strategy effectiveness

Coverage (%)	Minimize illness			Minimize deaths		
	Optimal strategy illness attack rate ^a	Random vaccination illness attack rate	Overall effectiveness ^b	Optimal strategy death rate ^c	Random vaccination death rate	Overall effectiveness
0	0.299	0.299		4.440	4.440	
10	0.143	0.240	0.404	2.108	3.299	0.361
20	0.105	0.184	0.429	0.716	2.297	0.688
30	0.021	0.133	0.842	0.335	1.505	0.777
40	0.014	0.082	0.829	0.144	0.867	0.834
50	0.010	0.043	0.767	0.111	0.438	0.747
60	0.008	0.023	0.652	0.084	0.218	0.615
70	0.006	0.012	0.500	0.070	0.132	0.470
80	0.005	0.008	0.375	0.056	0.085	0.341
90	0.004	0.005	0.200	0.047	0.061	0.230
100	0.003	0.003		0.044	0.044	

^aWe run the simulator 100 times for each strategy and present the mean overall attack rates.

^bWe calculate the effectiveness of implementing the optimal vaccine distribution strategy compared to completely random mass vaccination in the case of an Asian-like influenza pandemic. Effectiveness = $1 - Index_{opt}/Index_{rand}$.

^cDeath rate is represented as deaths per 10 000 people.

Table 7
Number of simulations for convergence^a (Asian-like Influenza)

Coverage (%)	Genetic algorithm		Random mutation hill climbing	
	Minimize illness	Minimize death	Minimize illness	Minimize death
10	1000	1000	860	740
20	1000	2500	27460	29430
30	3500	4500	66390	54090
40	3500	5500	148500	162530
50	2500	4500	75880	82340
60	3500	4000	89220	54530
70	2500	3500	108040	76940
80	2500	4000	63790	94380
90	1500	2000	850	1430

^aThe number of simulations until the optimal solution is found. Additional simulations are required for convergence to be determined, however this is the same for both GA and RMHC.

The GA is said to have converged if the best individual of the current pool does not change for four consecutive generations. The RMHC is said to have converged if the current individual does not change for 100 consecutive iterations. The convergence requirements are thus equivalent in terms of number of individuals evaluated after the optimal individual has been found. We can be confident, from the results in Table 7, that the GA converges more quickly than RMHC in this stochastic modelling implementation based on the large difference in convergence rates between the two algorithms, even though we do not have replications under each scenario. It would be possible to run the algorithm many times for each coverage and baseline attack rate calibration and determine a mean and variance for the number of

simulations required for convergence. However this was impractical in this situation as some runs of the RMHC algorithm may take up to 90 hours to complete. Mitchell et al. (1994) give a good comparison of the convergence between GA and RMHC and gives examples where one outperforms the other. It seems that the type of optimization that we encounter here is best suited for GA.

4. Discussion

The optimal vaccine distribution is sensitive to the nature of the spread of the influenza agent, the objective for control, and also the amount of vaccine available. GA, as well as RMHC, in conjunction with a properly

calibrated simulation model would allow us to obtain an approximation of the optimal vaccine distribution given that we understand the behavior of the next pandemic agent and that we determine an objective for the control of the agent.

It is of interest to compare the optimal distribution of vaccine obtained in our analysis with the recommended policies of the Centers for Disease Control and Prevention (CDC). For Interpandemic influenza, the CDC recommends routine vaccination of all people aged ≥ 50 , people at risk for influenza-related complications, and people who could transmit to others at high risk (Harper et al., 2004), regardless of the transmission patterns of the currently circulating strains. The pandemic influenza policy has not been fully established. However, during the swine influenza pandemic scare of 1976, efforts were made to mass vaccinate as many people as possible, with adults being vaccinated first (Neustadt and Fineberg, 1983). The mass vaccination campaign was halted after 40 million adults were vaccinated, a small number of which appeared to have developed vaccine-related Guillain–Barré syndrome. In addition, there was no real evidence that the swine influenza pandemic would materialize. We suspect that the real operational CDC pandemic influenza vaccination policy could follow the previous policy of 1976, and thus, we use random mass vaccination available quantities as the basis for comparison with the optimal policy.

We have shown that the optimal vaccine distributions are highly effective, especially when compared to random mass vaccination. Implementation of the optimal vaccine distribution for Asian-like pandemic influenza was found to be 84% more effective than random vaccination when there was only enough vaccine for 30% of the entire vaccination and the objective was to minimize illness. This optimal vaccination strategy involved concentrating vaccine in children, with the leftover vaccine going to middle aged adults. In this situation, given a population of 280 million people, we would be able to prevent 31 million illnesses following the optimal vaccination strategy rather than random mass vaccination. In the case of an actual pandemic, it will be critical to isolate the viral strain quickly to construct a vaccine, to identify quickly the age-specific attack rate pattern and finally to implement an optimal vaccination control strategy.

To control pandemic influenza with vaccines, we must isolate the pandemic influenza strain quickly as it is unlikely that pre-manufactured vaccine will be effective (Fedson, 2003). We assume that early isolates of the new pandemic strain would provide the seed strains for making vaccine that would be antigenically well matched to the wild circulating strain. Thus, we assume that this vaccine would have similar efficacy to the current interpandemic vaccines that are generally well

matched to the circulating wild strains each year. However, it is possible the match between a hastily constructed vaccine and circulating pandemic strain may not be good, especially if the wild virus undergoes mutation after the first wave at the emergence site. A sensitivity analysis on the assumed vaccine efficacy is a subject for future work. Our simulation model and optimization problem does not consider the efficacy of the distributed vaccine against future strains of the virus, nor does it consider the efficacy of repeated annual vaccination (Smith, 1999). We only consider the spread of a single strain of influenza and a vaccine that is effective against that strain.

The community on which we based our simulations consists of 10 000 people. The population is constructed to represent a cross-section of a typical American community. It represents the social connections that are responsible for the transmission of influenza and is not meant to be taken literally as a disconnected population. The model averages over a great deal of the social structure of the actual population that we are attempting to model. However, we are interested in simulating the effects of vaccination on the age-specific, final illness attack rates, and not the intricate transient dynamics of influenza epidemic. Thus, we believe that the relatively simple social structure modelled is adequate for the purposes of the analysis carried out here. In addition, since the influenza season generally lasts for about 4 months, usually between December and April of each year, actual epidemics occur in subpopulations and regions of the country at different times. We have not attempted to model this pattern for the whole country. If we assume that the epidemics spread to virtually the whole country with relative uniformity by the end of the season, then the optimal vaccination distributions that we obtain should be generalizable to the whole US population of 281 million people (Longini et al., 2004).

Some concerns about use of GA as an optimization tool are the amount of computer time it takes for one run. The GA may take three hours to give a good approximation of an optimal vaccine distribution under our current simulation model. The computer running time of GA has an $O(n)$ relation to the running time of the stochastic simulator. Hence, the GA will converge more slowly at a linear rate with increasing complexity and running time of the stochastic epidemic simulator. However, as the number of control variables to be optimized grows, the convergence of the GA may take exponentially longer. We will learn about the applicability of GA to larger numbers of control variables in the setting of a stochastic simulator as we use it for further optimization problems. Further problems will involve more complex control strategies and parameter estimation such as estimation of transmission probabilities from an infected individual to a susceptible individual.

GA is well suited for problems that have a complex fitness landscape as the recombination methods of GA are able to move the population out of local minima that a gradient search or hill climbing method may not be able to. The GA may also converge much more quickly than the RMHC algorithm because of the non-directional mutations employed by RMHC. In RMHC, we repeatedly randomly modify the currently best individual in hopes that its offspring are better than the original, without inducing any directionality into the random mutations. The GA attempts to breed two relatively fit individuals allowing the possibility of extracting the best genes from both individuals to create a more fit child. This approach takes advantage of genetic diversity that, as it seems, may be necessary to perform optimizations of this kind.

Although the optimal vaccination distributions found here are similar to those found by Longini et al. (1978) some years ago, the simulation model and optimization methods that we employ here are more general and comprehensive than those earlier efforts. When the next pandemic strain of influenza is identified somewhere in the world, vaccine formulation and production should proceed as quickly as possible. Once the age-specific illness attack rate patterns are identified, the epidemic simulation model with the current US population structure can be calibrated. Then the optimization model can be used to investigate the best vaccine distribution given the quantities available. Influenza antiviral agents could be used to slow transmission until vaccine is available (Longini et al., 2004).

Acknowledgment

This work was partially supported by grant R01-AI32042 from the National Institute of Allergy and Infectious Diseases.

References

- Anderson, R., May, R., 1991. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, New York.
- Belshe, R.B., Mendelman, P.M., Treanor, J., 1998. The efficacy of live attenuated, cold-adapted, trivalent, intranasal influenza virus vaccine in children. *N. Engl. J. Med.* 20, 1405–1412.
- Belshe, R.B., Gruber, W.C., Mendelman, P.M., 2000a. Correlates of immune protection induced by live attenuated, cold-adapted, trivalent, intranasal influenza virus vaccine in children. *J. Infect. Dis.* 181, 1133–1137.
- Belshe, R.B., Gruber, W.C., Mendelman, P.M., 2000b. Efficacy of vaccination with live attenuated, cold-adapted, trivalent, intranasal influenza virus vaccine against a variant (A/Sydney) not contained in the vaccine. *J. Pediatr.* 136, 168–175.
- Census Bureau, United States, 2001. *Census 2000*; <http://www.census.gov/>.
- Davis, L.E., Caldwell, G.C., Lynch, R.E., Bailey, R.E., 1970. Hong Kong influenza: the epidemiologic features of a high school family study analyzed and compared with a similar study during the 1957 Asian influenza epidemic. *Amer. J. Epidemiol.* 92, 240–247.
- Elveback, L.R., Fox, J.P., Ackerman, E., 1976. An influenza simulation model for immunization studies. *Amer. J. Epidemiol.* 103, 152–165.
- Fedson, D., 2003. Pandemic influenza and the global vaccine supply. *Clin. Infect. Dis.* 36, 1552–1561.
- Forrest, S., Mitchell, M., 1993. Foundations of genetic algorithms 2. In: Whitley, D. (Ed.), *Relative Building-Block Fitness and the Building-Block Hypothesis*. Morgan Kaufmann, San Mateo, CA, pp. 109–126.
- Greenhalgh, D., 1986. Control of an epidemic spreading in a heterogeneously mixing population. *Math. Biosci.* 80, 23–45.
- Halloran, M.E., Longini, I.M., Struchiner, C.J., 1999. Design and interpretation of vaccine field studies. In: Monto, A.S., Thacker, S.B. (Eds.), *Epidemiologic Reviews: Vaccines*, vol. 21, pp. 73–88.
- Halloran, M.E., Longini, I.M., Cowart, D.M., Nizam, A., 2002a. Community trials of vaccination and the epidemic prevention potential. *Vaccine* 20, 3254–3262.
- Halloran, M.E., Longini, I.M., Nizam, A., Yang, Y., 2002b. Containing bioterrorist smallpox. *Science* 298, 1428–1432.
- Harper, S.A., Fukuda, K., Uyeki, T., Cox, N., Bridges, C., 2004. Prevention and control of influenza: recommendations of the advisory committee on immunization practices. *MMWR* 53, 1–40.
- Hethcote, H., Waltman, P., 1973. Optimal vaccine schedules in deterministic epidemic models. *Math. Biosci.* 18, 365–381.
- Hill, A.N., Longini, I.M., 2003. The critical vaccination fraction for heterogeneous epidemic models. *Math. Biosci.* 181, 85–106.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor.
- Jordan, W.S., 1961. Mechanisms of spread of Asian Influenza. *Amer. Rev. Res. Dis.* 83, 29–35.
- Kavet, J., 1972. *Influenza and public policy*. Ph.D. thesis, Harvard University.
- Kiefer, J., Wolfowitz, J., 1952. Stochastic estimation of the maximum regression function. *Ann. Math. Statist.* 23, 462–466.
- Kilbourne, E.D., 1975. *The Influenza Viruses and Influenza*. Academic Press, New York.
- Longini, I.M., Ackerman, E., Elveback, L.R., 1978. An optimization model for influenza A epidemics. *Math. Biosci.* 38, 141–157.
- Longini, I.M., Halloran, M.E., Nizam, A., Wolff, M., Mendelman, P.M., Fast, P., Belshe, R.B., 2000. Estimation of the efficacy of live, attenuated influenza vaccine from a two-year, multi-center vaccine trial: implications for influenza epidemic control. *Vaccine* 18, 1902–1909.
- Longini, I.M., Halloran, M.E., Nizam, A., Yang, Y., 2004. Containing pandemic influenza with antiviral agents. *Amer. J. Epidemiol.* 159, 623–633.
- Mitchell, M., Holland, J., Forrest, S., 1994. When will a genetic algorithm outperform hill climbing? In: Cowan, J., Tesauro, G., Alspector, J. (Eds.), *Advances in Neural Information Processing Systems*. Morgan Kaufmann, San Mateo, CA.
- Neustadt, R.E., Fineberg, H.V., 1983. *The Epidemic that Never Was*. Vintage Books.
- Orvosh, D., Davis, L., 1993. Shall we repair? Genetic algorithms, combinatorial optimization, and feasibility constraints. In: Forrest, S. (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms*. Morgan Kaufman, San Mateo, CA, p. 650.
- Robbins, H., Munro, S., 1951. A stochastic approximation method. *Ann. Math. Statist.* 22, 400–407.
- Sharrar, R.G., 1969. National influenza experience in the USA, 1968–1969. *Bull. World Health Organ.* 41, 361–366.

- Smith, D., 1999. Variable efficacy of repeated annual influenza vaccination. *Proc. Natl. Acad. Sci.* 96, 14001–14006.
- Spears, W., De Jong, K., 1990. An analysis of multi-point crossover. *Proceedings of the Foundations of Genetic Algorithms Workshop*.
- Thompson, W.W., Shay, D.K., Weintraub, W., 2003. Mortality associated with influenza and respiratory syncytial virus in the United States. *JAMA* 289, 179–186.
- Weisstein, E., 2004. Robbins–munro stochastic approximation. *MathWorld*; <http://mathworld.wolfram.com/Robbins–MonroStochasticApproximation.html>.
- WHO, 2004. World Health Organization: Influenza; [http://www.who.int/mediacentre/factsheets/2003/fs211/en/](http://www.who.int/mediacentre/factsheets/fs211/en/).
- Wickwire, K.H., Guest, D., 1976. Optimal control policies for reducing the maximum source of a closed epidemic. *Math. Biosci.* 32, 1–14.