A Model for the Hanta Virus

Daniel Xiang
Tyler Devlin
Daniel Kunin

November 7, 2015

Introduction

Strategies combatting the recent outbreaks of infectious diseases, such as the 2014 Ebola outbreak in Western Africa or the 2009 outbreak of Avian Influenza, have demonstrated the effectiveness of mathematical models as a tool to mitigate fatalities and help determine the optimal allocation of resources to treat the largest number of patients. The World Health Organization is taking the current outbreak of the Hanta virus very seriously. After having learned from previous outbreaks of infectious diseases, we have implemented a mathematical model of both infection and treatment to inform our strategy in moving forward.

We started by modeling the Hanta virus in its human host from viral contraction until host death. Using information known about the virility of the virus and how the human immune system reacts, we were able to determine how long it would take before a patient starts exhibiting symptoms and how long they could live with the virus. We then modeled antibiotic treatment of a patient to determine how long they have to get medical help. By combining both of these models we were able to create a hypothetical timeline for a patient infected with the Hanta virus. This timeline will allow us to assess patients as they are reported to determine how far along the virus has spread and what options they have for treatment.

Modeling will only be part of the solution to the current outbreak of the Hanta virus. However, we are certain that with our model, we will be able to effectively direct medical resources and reduce the total loss of life.

Modeling Infection

Let $V(t)$ denote the population of the Hanta virus in the body of an infected person at time t where t is in hours. Here we make our first assumption.

Assumption: At time $t = 0$, there is exactly 1 copy of the virus in the body.

In other words, whenever someone contracts the virus, they contract only 1 copy of the virus, as opposed to several, which may be more realistic.

Although viral replication is a discrete process, we can model the viral population with a continuous function. Specifically, since the population doubles every hour when $V(t) < 10^6$, we have

$$\begin{aligned} V(t) &= 2^t \\ \Rightarrow \frac{dV}{dt} &= V \ln 2 \end{aligned}$$

If we plug in 10^6 for $V(t)$ we can solve for the time at which the viral population reaches one million. This time is $\log_2 10^6 = 19.93$ hours.

After 19.93 hours, the immune response begins to kick in. At maximum effectiveness, we are given that the immune system can destroy 200,000 viral copies per hour. We make the following assumption.

Assumption: As time increases, the effectiveness of the immune response also increases.

In other words, once the virus is detected, the strength of the immune response gradually increases, eventually leveling off at the maximum effectiveness.

Hence, we want to find a function $f(t)$ as a coefficient to our -200,000 term in our derivative that takes the value 0 at $t = 19.93$ and that approaches 1 as $t \to \infty$.

First we introduce a function $\sigma(t)$ such that $\sigma(t) \to 1$ as $t \to \infty$. The following sigmoidal function has the desired property,

$$\sigma(t) \doteq \frac{1}{1 + e^{-\alpha(t-19.93)}} \qquad \text{where } \alpha \in \mathbb{R}$$

Observe,

$$\lim_{t\to\infty} \sigma(t) = \lim_{t\to\infty} \frac{1}{1 + e^{-\alpha(t-19.93)}} = \frac{1}{1+0} = 1.$$

Notice that if we plug in $t = 19.93$, we find $\sigma(19.93) = \frac{1}{1+e^0} = \frac{1}{2}$, but we want our function to be 0 at the initiation of immune response. To avoid "jumps" in our immune response function, we construct the following $f(t)$.

$$f(t) \doteq \sigma(t) - \frac{1}{2}e^{-(t-19.93)}$$

So now, at $t = 19.93$, we have $f(19.93) = \frac{1}{2} - \frac{1}{2} = 0$.

As $t \to \infty$, the second term vanishes and the whole function $f(t) \to 1$, as desired.

This is consistent with reality because at $t = 19.93$ hours, our immune system kicks in. But the immune response requires time to reach its maximum effectiveness. So at $t = 19.93$ hours, the contribution should be 0, and as time progresses, the contribution should approach maximum effectiveness at -200,000 viral copies per hour.

We make the following modification to our model,

$$\frac{dV}{dt} = \begin{cases} V \ln 2 & \text{if } t \leq 19.93 \\ V \ln 1.5 - 200,000\left(\frac{1}{1+e^{-\alpha(t-19.93)}} - \frac{1}{2}e^{-(t-19.93)}\right) & \text{if } t > 19.93 \end{cases}$$

At $t = 19.93$ our derivative has no negative contribution, but as t increases, the coefficient function of the negative term approaches 1, yielding the maximally effective immune response in the overall derivative. The $\ln 1.5$ represents the switch from 200% viral growth/hr to 150% viral growth/hr.

We have an explicit solution for $t \leq 19.93$, but for $t > 19.93$, we will turn to numerical methods to solve for our solution $V(t)$.

We included the α coefficient in our sigmoidal function because we'd like to tune $f(t)$ to match the literature data for how effective the immune response is at some time t.

According to a paper written by Kumar and Sharma on infectious diseases[3], the immune response becomes effective after 3 hours. Here we make another assumption.

Assumption: After 3 hours, the effectiveness of the immune response is around 95% of the maximum effectiveness.

We use this data point to find an α that puts our model in the same order of magnitude as what has been observed in research.

3 hours after our immune response begins, we are at time $t = 3 + 19.93 = 22.93$ hrs.

So, plugging in $t = 22.93$, we solve for our tuning parameter α.

$$\begin{aligned}
f(22.93) &= .95 \\
&= \frac{1}{1+e^{-\alpha(t-19.93)}} - \frac{1}{2}e^{-(t-19.93)}\Big|_{t=22.93} \\
&= \frac{1}{1+e^{-\alpha(22.93-19.93)}} - \frac{1}{2}e^{-(22.93-19.93)} \\
&= \frac{1}{1+e^{-\alpha 3}} - \frac{1}{2}e^{-3} \\
\Rightarrow \frac{1}{1+e^{-3\alpha}} &= .95 + \frac{1}{2}e^{-3} \\
\Rightarrow 1 &= (.95 + \frac{1}{2}e^{-3})(1+e^{-3\alpha}) \\
&= (.95 + \frac{1}{2}e^{-3}) + e^{-3\alpha}(.95 + \frac{1}{2}e^{-3}) \\
\Rightarrow e^{-3\alpha} &= \frac{1-(.95+\frac{1}{2}e^{-3})}{(.95+\frac{1}{2}e^{-3})} \\
\Rightarrow \ln e^{-3\alpha} &= \ln\Big(\frac{1-(.95+\frac{1}{2}e^{-3})}{(.95+\frac{1}{2}e^{-3})}\Big) \\
\Rightarrow -3\alpha &= \ln\Big(\frac{1-(.95+\frac{1}{2}e^{-3})}{(.95+\frac{1}{2}e^{-3})}\Big) \\
\Rightarrow \alpha &= -\frac{1}{3}\ln\Big(\frac{1-(.95+\frac{1}{2}e^{-3})}{(.95+\frac{1}{2}e^{-3})}\Big) \\
&\approx 1.2197
\end{aligned}$$

So our model becomes

$$\frac{dV}{dt} = \begin{cases} V\ln 2 & \text{if } t \le 19.93 \\ V\ln 1.5 - 200,000\Big(\frac{1}{1+e^{-1.22(t-19.93)}} - \frac{1}{2}e^{-(t-19.93)}\Big) & \text{if } t > 19.93 \end{cases}$$

We use ode45 in Matlab to produce our solution $V(t)$ and use logical indexing in Matlab to find when the viral population is closest to 10^9. Below is a figure of our numerical simulation results of the second part of our piecewise derivative. The simulation begins with initial time and viral population of $(t_0, V_0) = (19.93, 10^6)$.

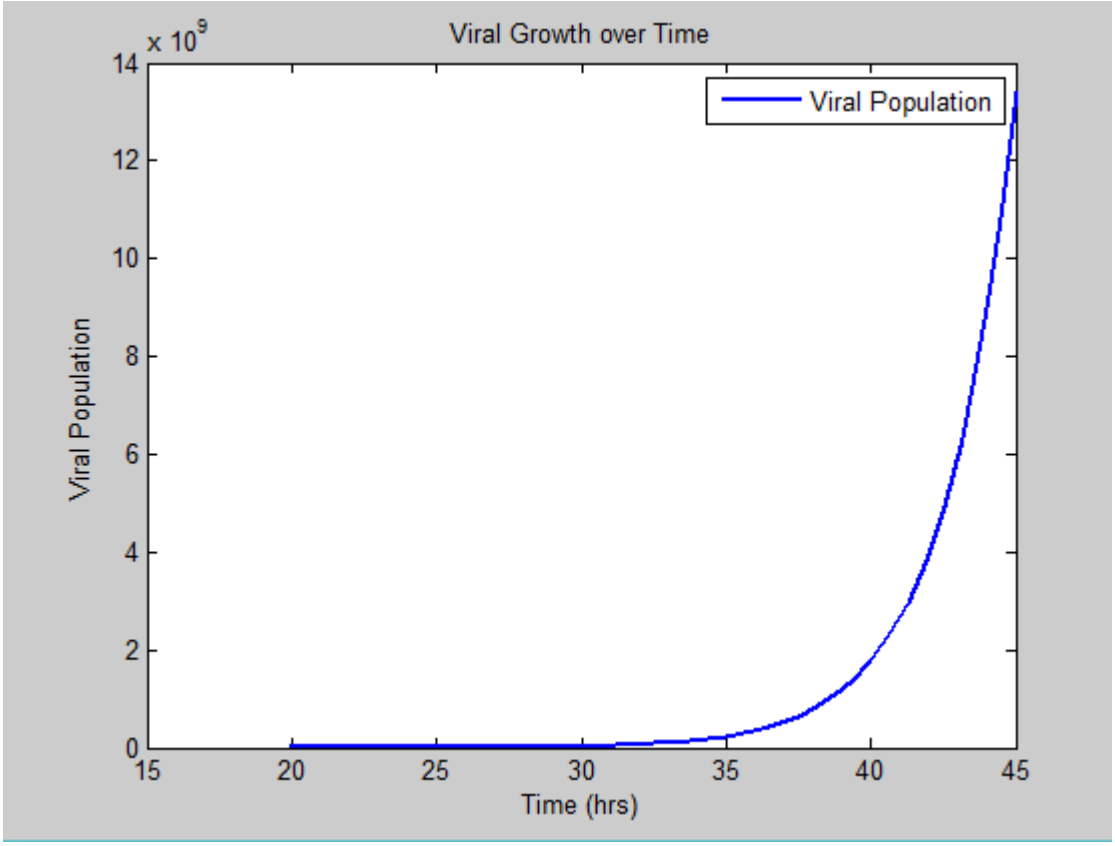

In particular, we take the square difference between every element in the solution vector $V(t)$ and 10^9, and take the min of this squared difference vector. This minimum corresponds to the viral population closest to one billion. The min function in Matlab also allows us to return the index of this population, so we can find the corresponding time in the t vector also produced by ode45 simulation.

This yields the time $t = 38.7742$ hours. Thus, if a patient doesn't get treatment before 38.77 hours after contracting the virus, she will be beyond saving.

Modeling Treatment

We are given that after antibiotics have been distributed, the immune system combined with the antiobitics kill 500 million copies of the virus per hour. So after antibiotics have been distributed, our model becomes

$$\frac{dV}{dt} = \begin{cases} V \ln 2 - 5(10^8) & \text{if } t \leq 19.93 \\ V \ln 1.5 - 5(10^8) & \text{if } t > 19.93 \end{cases}$$

We recognize the autonomous nature of this ODE and solve for the equilibrium. We only consider the case for $t > 19.93$ because one can simply observe that the steady state corresponding to $t \leq 19.93$, which is $V^* = \frac{5(10^8)}{\ln 2}$ is several orders of magnitudes greater than our solution can reach by time $t = 19.93$, which is around only 1 million.

So for $t > 19.93$, we have

$$\begin{aligned} 0 = \frac{dV}{dt} &= V^* \ln 1.5 - 5(10^8) \\ \Rightarrow V^* &= \frac{5(10^8)}{\ln 1.5} \approx 1.23(10^9) \end{aligned}$$

A simple sign analysis shows that this is an unstable equilibrium, meaning that if the population is below this level, the derivative will be negative, and the population will diminish to 0 as $t \to \infty$. Conversely, when the population is above this level, the derivative will be positive, and the population will diverge to $+\infty$ as $t \to \infty$.

This analysis is consistent with our given information because we are told that after the viral population exceeds 1 billion, the patient is incurable, and our qualititative analysis says that once the population is above 1.23 billion, the viral population diverges to $+\infty$.

We also verified this with simulation. See the figures below.

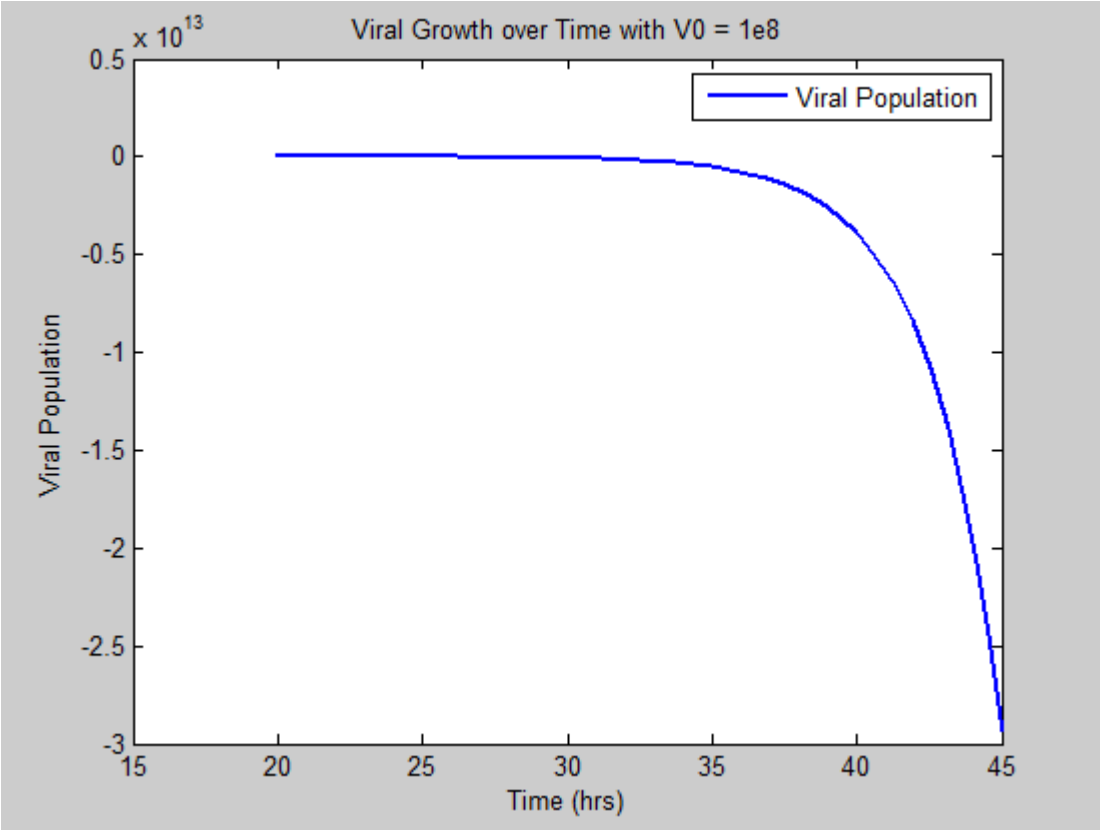
Viral Growth over Time with V0 = 1e8
x 10^13
Viral Population
Viral Population
Time (hrs)
0.5
0
-0.5
-1
-1.5
-2
-2.5
-3
15
20
25
30
35
40
45

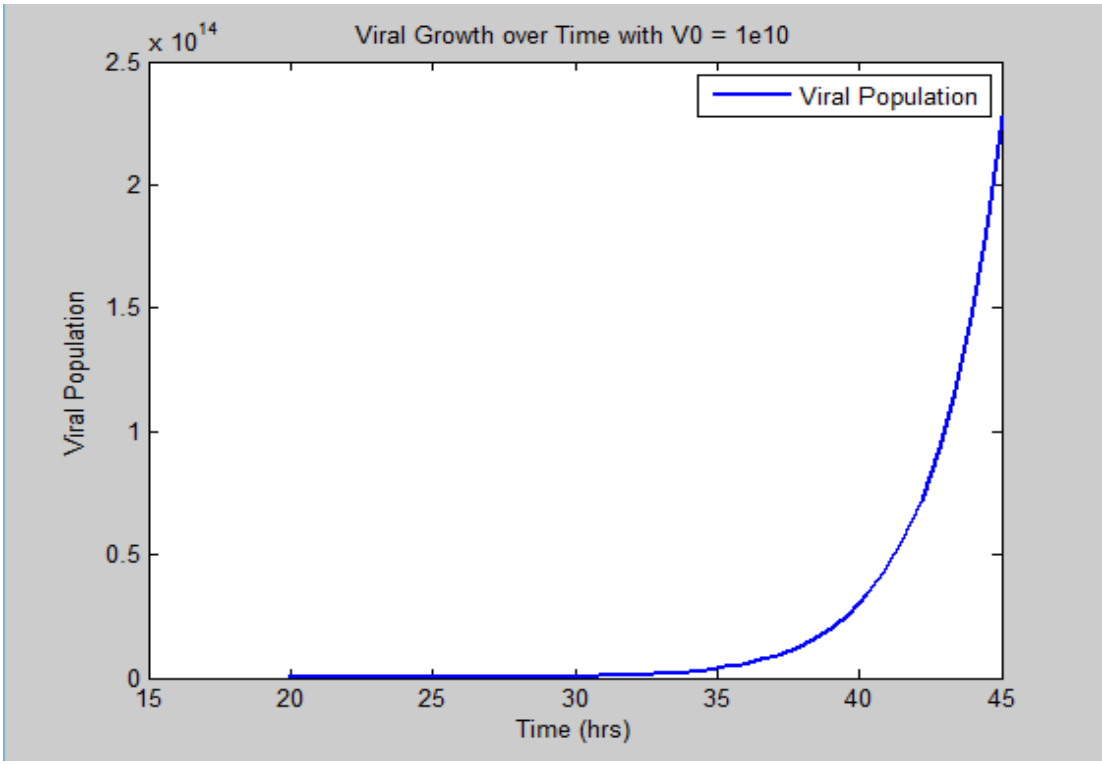
Viral Growth over Time with V0 = 1e10
x 10^14
Viral Population
Viral Population
Time (hrs)
2.5
2
1.5
1
0.5
0
15
20
25
30
35
40
45

In the first figure, we didn't impose the lower population bound of 0, so the simulation just showed the population diverging downard towards $-\infty$. In reality, it would just approach 0 and stay there.

In the second figure, we simulated with an initial condition bigger than the equilibrium, and saw that the population was diverging to $+\infty$. This is what we expected.

Since the unstable equilibrium we solved for analytically is bigger than 10^9, we can conclude that if the antibiotics are issued before the viral population reaches 10^9, then the viral population $V(t)$ will diverge downards until it hits zero, and the patient will be saved.

And finally we are given that a person who's viral population reaches 1 trillion copies will die. So we simulate and find the time corresponding to $V(t) = 10^{12}$ in a similar method described above. This yields the time that it takes for the virus to kill someone to be about 55.47 hours, assuming they weren't treated with antibiotics.

The Timeline

Now that we've developed some basic models for the viral population with and without antibiotic treatment, we can construct a timeline for a patient infected with the Hanta virus. Using this timeline, we can prioritize the treatment of patients depending on their position on the timeline: patients closer to the point of no return should be prioritized. Those beyond that point will receive comfort in their inevitable passing.

A patient begins feeling the full symptoms due to their immune response at around time $t = 22.93$ hours as shown in the previous section. In order to calculate when the patients became infected, we add the amount of time that has passed since they first started to experience symptoms to 22.93 hours.

For example, if a patient reports that she first experienced symptoms 24 hours ago, we would estimate that she contracted the virus $24 + 22.93 = 46.93$ hours ago, and would place her at $t = 46.93$ hours on the timeline.

However, we can't assume that the amount of time the patient reports is completely accurate. It could be the case that they assumed they only had a bad cold or flu, and as a result reported that they didn't start to experience the Hanta symptoms until later. Also, since a Hanta outbreak has been announced, people may be likely to report symptoms earlier. We make the following assumption.

Assumption: The patient's estimate for the amount of time that has passed since first experiencing symptoms is normally distributed among patients.

We give a 95% confidence interval for ± 3 hours of what the patients report, meaning that we are 95% confident that the true amount of time that has passed is within ± 3 hours of what the patients report.

We can now represent each patient as an interval of length 6 on the timeline. This interval is centered around the prediction based on their estimate of the time that has passed since first experiencing symptoms. The right bound of this interval should not be allowed to pass the point of no return if the patient is to be saved.

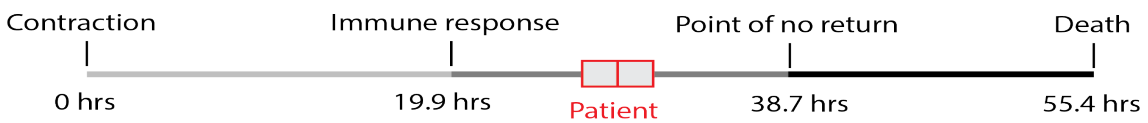

Summary

In summary, it takes 19.93 hours for the immune response to begin once the virus has been contracted. Once an individual has contracted the virus, they have 38.77 hours to get help from medical authorities.

To determine when patients became infected, we add their estimate of time passed since first experiencing symptoms to 22.93 hours. We then decide on how quickly they need to be treated by representing them as an interval on the timeline.

In order to determine how much time we have to treat the patient, we subtract the right end of their interval from 38.77 hours, the point of no return.

For example, if a person says they first started experiencing symptoms 12 hours ago, we compute their interval endpoints by adding and subtracting 3 to the sum of their prediction and 22.93 hours, yielding the interval

$$(22.93 + 12 - 3, 22.93 + 12 + 3) = (31.93, 37.93)$$

We then subtract the right bound of this interval from the point of no return, $t = 38.77$ hours, giving us the amount of time we have to treat the patient in order to be 95% confident they will be saved, which is $38.77 - 37.93 = 0.84$ hours. We'd better hurry!

When approached by more than one patient, we implement the timeline described above, representing each patient as an interval whose endpoints depend on their estimated time since first experiencing symptoms. In order to allocate resources efficiently, those patients who's right bounds are closer to the point of no return are treated before the patients who are further left on the timeline.

One of the virtues of our model is its great versatility. We can simply change a couple parameters in order to effectively model any exponential viral population growth. This model provides a basis on which we can improve, tune, and apply to other problems involving exponential growth followed by damping after a certain threshold, which constitute a large class of problems limited not only to population growth – we can take similar approaches to model diverse problems such as radioactive decay, credit payments, and temperature cooling/heating to name a few.

The weaknesses of our model lie in the unavoidable nature of making simplifying assumptions. Our first assumption was to say that every contraction of the virus involves contracting only one copy of the virus, where in fact it is more likely that people will contract more than one copy of the virus at a time. For example, contracting two copies at once would cause the times calculated throughout the report to be shifted to the left by an hour, since we start with double of what we started with when only contracting one copy of the virus.

Our second assumption is a reasonable one, without many downfalls. As time increases, the immune response becomes more effective. This is generally true.

Our third assumption generates a certain margin of error. The paper cited didn't give a precise number for the effectiveness of the immune response after 3 hours, so a reasonable estimate was made. Along with any estimate comes a certain margin of error.

Our fourth assumption has similar downfalls to the third. The idea was to provide some margin of error in the average individual's ability to estimate a time period so that we could devise a strategy to allocate our limited resources assuming we can't just treat everybody at once.

In conclusion, our mathematical model is far from perfect, but it at least provides a basis through which we are able to observe valuable insights into the nature of the viral population growth problem. We can then use these insights to inform our colleagues in devising effective strategies in order to combat the Hanta virus.

References

[1] Boyce, W., and DiPrima, R. Elementary Differential Equations and Boundary Value Problems 2d ed. New York: Wiley, 1969. Print.

[2] Foley, H. "Time Perception." Web. 8 Nov. 2015

[3] Kumar, V., and A. Sharma. "Neutrophilis: Cinderella Of Innate Immune System." International Immunopharmacology: 1325-334. Print.

[4] Vardeman, S. "What about the Other Intervals?" The American Statistician 46 (3): 193-197.