A model for zombie infestations at Wesleyan

Dave Constantine

September 25, 2017

1 Introduction

In this project we develop and investigate a differential equations model for a zombie infestation of Wesleyan's campus. We came up with several models which might describe the spread of zombies across campus and we will describe below how these models were arrived at, and how they make some intuitive sense. Using some limited data about the infestation, we used one of the models to predict what will happen in this infestation using a numerical method, and we report on the results below. Then we discussed a possible alteration of the model to account for active killing of the zombies and we used it to figure out how many zombies per hour must be killed if the outbreak is to be stopped.

2 The models

We were asked to model the following situation. A zombie outbreak begins at Wesleyan with one zombie at 12am, which increases to 3 by 1am. By 6am there are 20 zombies and the original zombie died at 4am. We are told that we expect zombies to die between 4 and 24 hours after infection. We also know that there are 4000 people on campus, and that this number is fixed, as campus has been sealed off. This data is summarized below:

time (hr)	Humans	Zombies	Dead
0	3999	1	0
1	3997	3	0
2			0
3			0
4			1
5			1
6	3979	20	1

In this table, time refers to hours since midnight. We are not told exact zombie populations between 2am and 6am.

To model the outbreak we introduce the following notation: $H(t)$ will be the number of humans at time t, $Z(t)$ the number of zombies, and $D(t)$ the number of dead zombies. We know from the description of the situation that the total number of individuals is always 4000 or

$$H(t) + Z(t) + D(t) = 4000.$$

We will assume that the outbreak takes place over a relatively short period of time, so zombies are the only cause of death that has any real chance of happening.

Since we do not have full information about the way the populations evolve, we attempt to build a model by describing the rate of change of each population in a reasonable way. There are two sources of population change: humans becoming zombies, and zombies dying. The first happens more often when there are more zombies and more humans, and we know that it would drop to nothing if there were either no zombies present, or no humans left to infect. Without other information we decided to model this rate of infection by $\alpha H(t)Z(t)$ since it is a relatively simple expression that matches the three criteria mentioned above. The parameter α will be estimated below. It is positive and the larger α is, the faster the zombie population increases.

The rate of zombie death is trickier to estimate. One possibility is to suppose that all zombies die 14 hours since they were created, since this is the midpoint of the 4-24 hour lifespan they have. In this case, the rate of zombie death at time t is the rate of zombie creation 14 hours earlier, at time $t-14$, or $\alpha H(t-14)Z(t-14)$. Another possibility is to suppose that a fixed proportion of the zombies die in any given hour. In this case the rate of death is $\beta Z(t)$ where β is a parameter between 0 and 1. For this paper, we decided to try the first model.

Using these assumptions, we get the following DE model for the human, zombie, and dead populations:

$$\begin{aligned}
\frac{dH}{dt} &= -\alpha H(t)Z(t) \\
\frac{dZ}{dt} &= \alpha H(t)Z(t) - \alpha H(t-14)Z(t-14) \\
\frac{dD}{dt} &= \alpha H(t-14)Z(t-14).
\end{aligned}$$

One thing to note about this model is that $\frac{dH}{dt} + \frac{dZ}{dt} + \frac{dD}{dt} = 0$. This coheres with the fact that the total number of individuals stays constant.

3 Qualitative Analysis

We can draw a few qualitative conclusions from this model. First, there are equilibrium solutions when $Z(t)$ is always equal to zero, or when $H(t)$ and $H(t-14)$ are both equal to zero. In the first, no outbreak occurs at all. This makes sense as an equilibrium, although it is not what happens in our situation. In the second, the outbreak has completely killed off all the humans for more than 14 hours, and now the last of the zombies have died off as well – clearly an equilibrium situation.

We can also note briefly that $\frac{dH}{dt}$ is always negative or zero, so the human population always decreases, and $\frac{dD}{dt}$ is always positive or zero, so the dead population always increases. Because of the $t-14$ factor, the system is not autonomous, so we could not draw a phase portrait.

4 Numerical solution

Since we don't know how to write down an exact solution for this model, we used the following numerical method. Assuming we know the populations as time t, we can approximate their value at time $t+1$ using the rule:

$$Pop(t+1) = Pop(t) + \frac{dPop}{dt}\Delta t.$$

This is Euler's method with a step size of $\Delta t = 1$. For our particular system, it yields:

$$\begin{aligned}
H(t+1) =& H(t) - \alpha H(t)Z(t) \\
Z(t+1) =& Z(t) + \alpha H(t)Z(t) - \alpha H(t-14)Z(t-14) \\
D(t+1) =& D(t) + \alpha H(t-14)Z(t-14).
\end{aligned}$$

Using the initial data in the table above, we can then find approximate solutions to this system of DEs using Euler's method in an iterative way. Since we need to use times 14 hours in the past, we filled in 0 for $D(t)$ before midnight, and 1 for $Z(t)$ stretching back to 14 hours before the first zombie died. We used an Excel spreadsheet to carry out this calculation. A portion of the calculation is below. Note that in Excel, we formatted the data to round populations off to the nearest integer, but in reality, the simulation has fractional values for the population.

To carry out these calculations, we needed a value of the parameter α. Since not enough detail was given to come up with a sensible value *a priori*, we tried several values until the calculations matched the given data relatively closely. In particular, we looked for a value of α which would match the 20 zombie population at time $t = 6$. In the end, we settled on $\alpha = .00015$, which gave $Z(6) \doteq 17$. A graph of the human, zombie, and dead populations for the first 40 hours is below.

	A	B	C	D	F	G	H	J	K
1	t	H(t)	Z(t)	D(t)	dH/dt	dZ/dt	dD/dt	alpha=	0.000150
2	-14	4000	0	0					
3	-13	4000	0	0					
4	-12	4000	0	0					
5	-11	4000	0	0					
6	-10	4000	0	0					
7	-9	4000	0	0					
8	-8	3999	1	0					
9	-7	3999	1	0					
10	-6	3999	1	0					
11	-5	3999	1	0					
12	-4	3999	1	0					
13	-3	3999	1	0					
14	-2	3999	1	0					
15	-1	3999	1	0					
16	0	3999	1	0	-0.59985	0.59985	0		
17	1	3998	2	0	-0.9595261	0.95952607	0		
18	2	3997	3	0	-1.5346431	1.53464308	0		
19	3	3996	4	0	-2.4538973	2.45389734	0		
20	4	3993	7	0	-3.9223186	3.92231862	0		
21	5	3990	10	0	-6.2656972	6.2656972	0		
22	6	3983	17	0	-9.9995457	9.39969567	0.59985		
23	7	3973	26	1	-15.576565	14.9767145	0.59985		
24	8	3958	41	1	-24.406473	23.8066234	0.59985		
25	9	3933	65	2	-38.301685	37.701835	0.59985		
26	10	3895	103	2	-59.955892	59.356042	0.59985		
27	11	3835	162	3	-93.17776	92.5779097	0.59985		
28	12	3742	255	4	-142.87571	142.275858	0.59985		
29	13	3599	397	4	-214.22725	213.627396	0.59985		
30	14	3385	610	5	-309.93654	309.336687	0.59985		
31	15	3075	920	5	-424.22863	423.269101	0.95952607		
32	16	2651	1343	6	-533.98431	532.449665	1.53464308		
33	17	2117	1876	8	-595.455	593.001105	2.45389734		
34	18	1521	2469	10	-563.24296	559.320644	3.92231862		
35	19	958	3028	14	-435.05257	428.786878	6.2656972		
36	20	523	3457	21	-271.09069	261.091148	9.99954567		

Figure 1: A portion of our numerical calculation

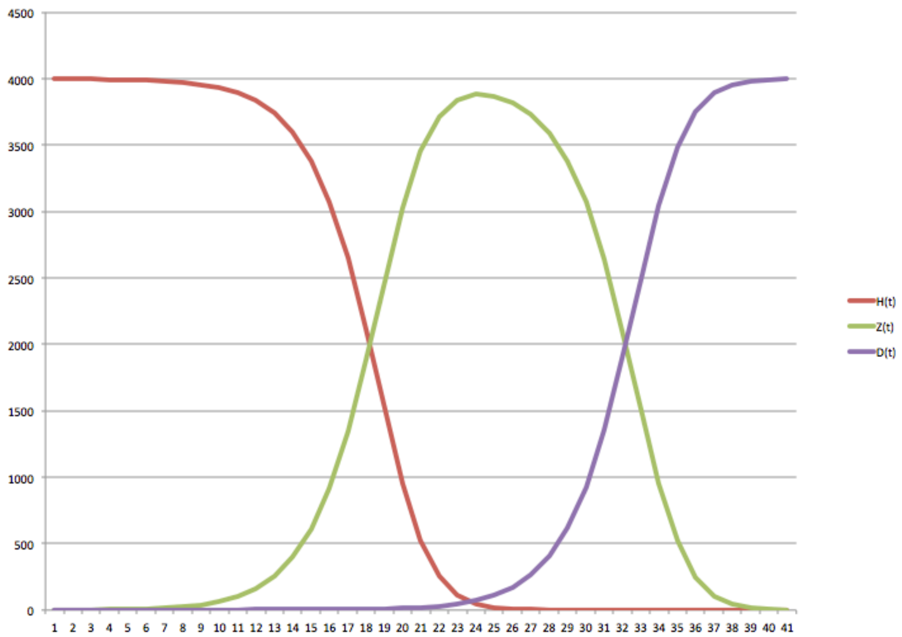

Figure 2: Results of our numerical simulation

5 A model with zombie-killing

Since we see from the graphs that this zombie outbreak will eventually kill us all, we investigated what could be done about this situation. One possibility is that, starting at 6am when we become aware of the outbreak, we send out a group of people to kill zombies. Let's assume that they kill γ zombies an hour. From that point on, we have a new set of differential equations:

$$\begin{aligned}\frac{dH}{dt} &= -\alpha H(t)Z(t) \\ \frac{dZ}{dt} &= \alpha H(t)Z(t) - \alpha H(t-14)Z(t-14) - \gamma \\ \frac{dD}{dt} &= \alpha H(t-14)Z(t-14) + \gamma.\end{aligned}$$

We are assuming here that hunting the zombies does not increase the danger of infection, i.e. does not increase α, and that no zombie hunters die directly from their efforts, rather than via infection. This model is also not perfect because the $-\alpha H(t-14)Z(t-14)$ factor for zombie death won't be accurate any more, since some of those zombies will have been killed already. If the zombie population is fairly large compared with γ, however, this is only a slight error.

We ran this new model through the same calculations, starting the killing at 10 hours, when there were about 100 zombies. The results with killing rates of 50 and 57 per hour are shown below.

6 Results

For our original model, we found that the zombies eventually wipe out all the humans. At the 29 hour mark, all humans are gone and by 40 hours, all the zombies have died off so that all individuals are dead. The curves graphed above show an outbreak that begins slowly, then very quickly infects all the humans, with almost all of the infections coming between 15 and 21 hours. This indicates that the outbreak could get out of control very quickly. By the time deaths start happening at any significant rate, the humans are mostly gone, so the last 16 hours or so of the simulation mainly consist of zombies dying off. The zombie population is roughly symmetric about its maximum at around 23 hours, although it does seem to die off a little bit slower than in spreads.

For the models with zombie killing, we found that it took a fairly large killing rate to make any difference. A killing rate of 50 per hour, roughly half of the active zombies when the zombie hunting started, clearly changes the results and decreases the speed of the outbreak and the total number of zombies at any time, but the human population still dies out. At 57 zombies

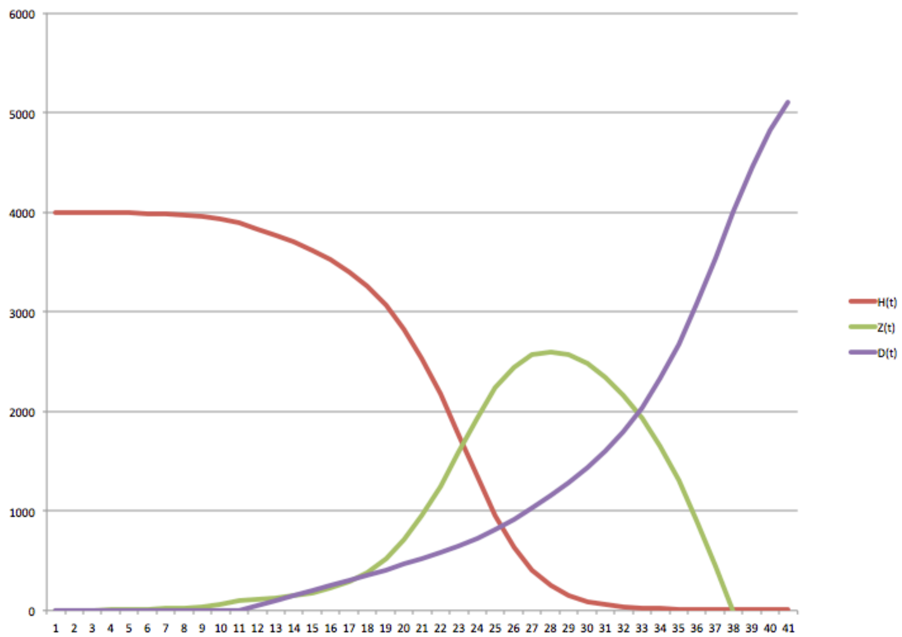

Figure 3: Results with killing rate of 50/hr

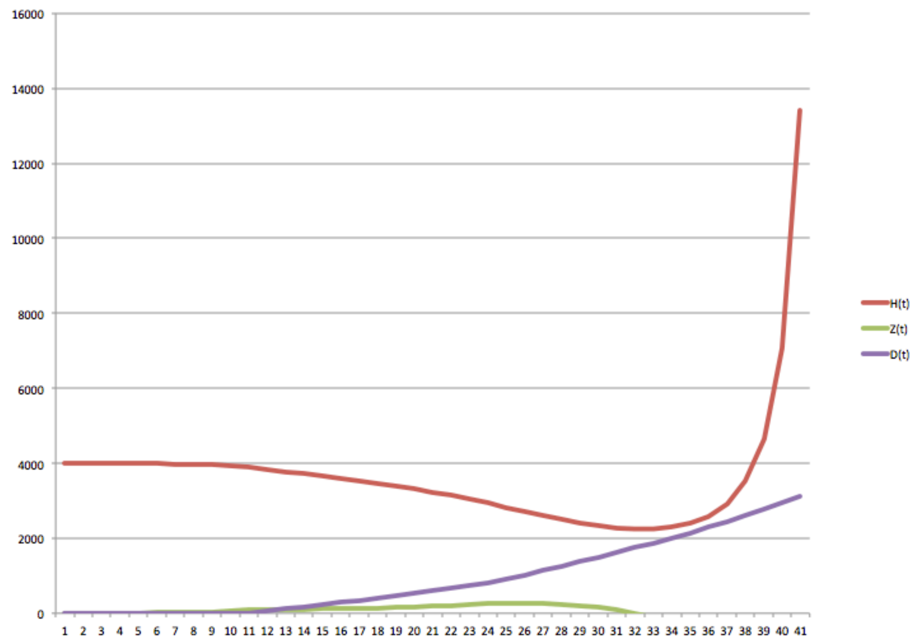

Figure 4: Results with killing rate of 57/hr

per hour, however, the outbreak is quickly contained, ending by 33 hours, with a loss of less than half of the human population.

7 Analysis

There are several possible sources of error in our original model. Some are in the original model itself. We supposed that all zombies die exactly 14 hours after infection, but the description we were given we just know that they die between 4 and 24 hours from infection. If the zombies might have different lifespans, we might have to modify the model. We also picked the simplest model that fit the data. With more information we might be able to try more detailed models with more parameters. For instance, we could suppose that the rate of zombie creation is $\alpha H(t)^{p_1} Z(t)^{p_2}$ for some $p_1, p_2 > 0$. This generalizes the model above where p_1 and p_2 were both 1. These parameters might enable us to fit more detailed data more precisely, and they would still fit the qualitative criteria for the model that we noted above.

The model with zombie killing seems to have larger problems. As noted above, it really only makes sense if the killing rate is not too large compared with the population. But we found that it took a fairly large killing rate (about half the population) before the zombie hunting made much of a change in the predictions, calling them into question somewhat. It's also worth noting that these models stop making sense once the number of zombies drops below the killing rate. We can see this in the solution graphs, where the number of dead exceeds 4000. This comes from the flaw of having 50 or so zombies 'killed' when fewer than 50 are actually still there. This killing model could be improved by adjusting the natural death rate to account for zombies already killed, and by making sure the killing rate drops when the number of zombies to kill is too small.

Another source of error is Euler's method itself. A step size of 1 is rather large. The model should certainly be run with smaller step sizes to get increased accuracy.

The model has several strengths, also. It fits our qualitative idea of what an outbreak would look like quite well, and it is possible to draw from it a sense of how quickly an outbreak could spread. It is also quite easy to work with and relatively simple, although it produces solutions which do not appear to be well-known functions. It is also fairly easy to adapt by adding new parameters.

In addition, it should be applicable to many other situations. For instance, zombies could be replaced by any other sort of infection. Another possible application would be to the spread of a rumor which, a certain time after hearing it, the hearer learns is not true.