Capturing the impact of patient variability in a novel cancer treatment using Bayesian inference

Mikaela Westlake

Supervised by Dr Adrianne Jenner & Dr Leah South

Queensland University of Technology

Contents

1   Abstract

Capturing patient outcomes in humans is a tricky task due to the variation in both an individual patients cancer growth and their response to treatment. However, in a set of identical nude mice, we are able to apply the same treatments and predict their outcomes. This is done with an oncolytic adenovirus which is a genetically engineered virus that only infects cancerous cells. This virus replicates until the cell undergoes lysis, cell death, releasing the virus back into the body for reinfection. This is already an approved treatment for Melanoma in the USA.

Data from a paper by Kim et al. (2011) on the growth of breast cancer tumours with and without was studied and modelled in a paper by Jenner et al. (2018). From this came an Ordinary Differential Equation model and a set of parameters fit for each mouse using least squares. The goal of this research was to take this previous work and produce a set of posterior distributions that better describe these parameters. This was done through the use of MATLAB 2021b and Bayesian inference.

The posteriors achieved produced a mean value, when compared to the estimates from Jenner et al. (2018), showed support as they were predominately within $\pm 20\%$. With new, more certain parameters values, new treatments could be run for each mouse to predict their outcomes. In all mice, increasing the frequency and duration of treatment produced a decrease in the rate of growth of the tumour. Increasing the amount of virions received at each dosage also decreased the overall tumour size but at a lower rate. Further work in this area includes investigation into new and more effective dosage regimes as well as comparing the results from breast cancer to other cancers such as liver and ovarian.

2   Statement of Authorship

The concept for this project is based on works in a previous paper by Dr Adrianne Jenner. This provided the Ordinary Differential equation model used heavily throughout this research as well as the prior parameter estimates. Dr Leah South over saw and advised in the Bayesian Inference process. Mikaela Westlake implemented the ODE model into MATLAB 2021b as well as the Bayesian inference process. Analysis of results was performed by Mikaela Westlake under the advice and guidance of Dr Adrianne Jenner and Dr Leah South.

AMSI

3   Introduction

The ability to predict and model patient outcomes in novel cancer treatments is a powerful tool. Cancer-targeting clinical trials generally fail due to the inability to distinguish between patients that respond to treatment and those that don't. There does exist an FDA-approved treatment for solid tumours such as melanoma, that is in use today (Tragner et al., 2020). This treatment is oncolytic virotherapy (OV) which involves the genetic modification of a virus, making it only infect tumour cells. OV typically uses an adenovirus which is genetically engineered to express particular genes that cause it to only infect cancerous tissue (Wares et al., 2015). Adenoviruses are typically used due to their ability to infect cells that both do and do not undergo cell division, but carry very mild symptoms with their infection. Some common adenoviruses include the common cold, gastroenteritis and pneumonia (Pitone, 2022). There are further modifications to the adenovirus apart from its tumour seeking ability that can further protect it from immune detection or alter its place of processing in the bloodstream (Kim et al., 2011).

The virus is effective as a cancer treatment due to its natural ability to infect, replicate. Figure 1 displays the cycle of infection seen in this treatment. First the virus is injected subcutaneously into the subject, this contains what is referred to later as the dose. This dose of virus cells are then able to infect any susceptible tumour cells. Once infected, the virus begins replicating. This continues until the tumour cell reaches its carrying capacity and undergoes lysis. This kills the cancerous cell and releases a portion of the virus back into the tumour for reinfection.
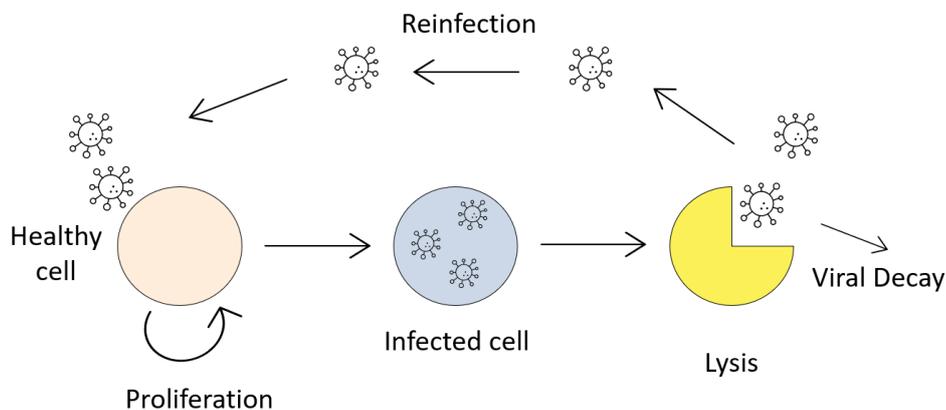

Figure 1: Schematic of infection and cell cycle with OV

A paper by Kim et al. (2011), studied the growth of susceptible breast cancer cells in genetically identical nude mice. These mice were bred to have no functioning immune system, eliminating a potential inhibiting factor of this treatment. In this study there were 4 different treatment types observed; A control, a "nude" Adenovirus, a PEG modified Adenovirus and a PEG-HER Adenovirus. The latter two were further modifications to the virus to increase its length of life in the blood stream by shielding it from the immune system. The control

sample set allowed continuous uninterrupted growth of the breast cancer tumour for 60 days. The adenovirus treatment was a single dose of the virus (10^{10} cells) administered on day 0, 2 and 4. A paper by Jenner et al. (2018) analysed the results achieved by Kim et al. (2011) and attempted to see if there was a model or a particular set of parameters that could describe the growth of this tumour and then model the effects of the different treatments. This was done by establishing an ordinary differential equations model (ODE) and using a least squares fit in MATLAB for a set of parameters for each mouse. This method only assigns one estimated value to each of the parameters.

The aim of this research project was to utilise the previous findings and ODE model from Jenner et al. (2018) and apply Bayesian inference to see if we can generate a distribution of parameters for each mouse. This would aid in future endeavours to form a distribution of potential parameters that could be assigned to a patient with particular predispositions or complications such as age and medical history. This would allow us to simulate their treatment outcomes prior to the commencement of treatment.

4 Background

4.1 The ODE Model

The ODE model outlined in equations $1 - 4$ established by Jenner et al. (2018). Each equation describes the rate of change in population of a cell type with respect to time. The solutions to these equations are the populations of each cell type at time t. Below in Table 1 is a description of the parameters in the ODE model. There are two rates of decay of populations as d_V and d_I as well as the rate of proliferation/growth of the tumour cell population r. α describes the amount of viral cells released from a lysed cell which occurs at the rate d_I. Finally, the rate at which the virus infects tumour cells is described by β, and the amount of cells received per injection of treatment is V_0.

Table 1 - Parameters in Ordinary Differential Equation Model

Parameter	Units	Description
d_V	day^{-1}	Viral Decay Rate
α	virus cells $\times 10^9$	Viral Burst Size
d_I	day^{-1}	Lysis Rate
r	day^{-1}	Tumour Proliferation Rate
β	day^{-1}	Virus Infection Rate
V_0	virus cells $\times 10^9$	Treatment Viral Dose
L	$cells \times 10^6$	Carrying Capacity

Using the above parameters, the set of differential equations were established and are described below. Equation 1 represents the rate of change in the population of virus cells with respect to time. This begins with the number of virus cells injected into the system at any time point where u_V is described in equation 4. This influx of viral cells then has the number of cells that have naturally decayed subtracted. This is the term $d_V \times V$ where V is the solution to the ODE. Finally, this population is increased by the number of viral particles released from the lysed infected cells.

Equation 2 represents the change over time in the population of susceptible tumour cells. This population grows according to a Gompertz growth function. Gompertz growth functions are typically used in modelling biological populations due to the models ability to capture the plateau when its reaches its carrying capacity (Tjørve et al., 2017). This term takes the proliferation rate r and multiplies it by natural log of carrying capacity, L, divide the number of susceptible cells, S, times the current population S. In equation 2, S is our solution variable. From this growth, we also subtract the number of cells infected by the virus represented by the rate of infection β times the current population of susceptible cells S, and the virus population V. This term is also divided by T, which is the total number of cells in the tumour being $T = S + I$.

The final population equation 3 is the rate of change of the infected cell population. This begins with the same susceptible to infected cell conversion term as previously described. It also has the rate of lysis of the infected cells subtracted from it. This is dependent of the rate at which they undergo lysis, d_I, times the current population size I, also being the solution to this ODE. Equation 4 of this system represents how the treatment is administered. This is done by a delta function and allows treatment to be given on day 0, 2 and 4 with an amount of cells specified by the value of V_0.

The full system from Jenner et al (2018) can be seen below

$$\frac{dV}{dt} = u_v(t) - d_v V + \alpha d_I I \tag{1}$$

$$\frac{dS}{dt} = r log\left(\frac{L}{S}\right) S - \frac{\beta SV}{T} \tag{2}$$

$$\frac{dI}{dt} = \frac{\beta SV}{T} - d_I I \tag{3}$$

$$u_V(t) = V_0(\delta(t) + \delta(t-2) + \delta(t-4)) \tag{4}$$

This ODE model was implemented into MATLAB (2021b) and solved using the inbuilt function $ode45$. Each of the solutions were output under the matrix y, where each column represents a different cell type, and each row represents their population at corresponding time t.

4.2   Bayesian Inference

Bayesian Inference is a statistics tool which allows us to gain information about our parameters before ever seeing them. At its core, Bayesian inference is an extension of Bayes theorem which is outlined as in figure 2. The probability of our data occurring, y, given a parameter, for example X has occurred, is proportional to the probability of X occurring given the data exists, times the probability of X (Hrouda-Rasmussen, 2021). The value of X here could be any of the paramters we chose to sample. in this case we will use r, L, β and σ. Here, $P(X)$ gives us our prior, which encompasses our prior beliefs about the parameters before seeing the data. These priors are to be based on Dr Jenner's beliefs about the parameters from prior work. Our posterior achieved will be in the form of a distribution which will tell us the most plausible values of X.

$$P(X|y) \propto P(y|X)P(X)$$

Posterior   Likelihood   Prior

Figure 2: Bayes Law as a Proportionality

4.3   Method

A Bayesian inference technique known as importance sampling from the prior was used to estimate the posterior distribution. To perform this technique, the previously estimated values were used to form a general range for three parameters r, L, β, and σ. From this range, a uniform distribution was formed based on these values and sampled N number of times. A standard deviation matrix was generated using values sampled from a gamma function. Using these distributions, the N samples from each parameter were stored as columns in a matrix where each row represents a different combination of parameters. Each set was run through the ODE in MATLAB 2021b using the *ode45* function. In the ODE contained 3 separate executions of *ode45*. Each call of it ran for a set amount of time then took the last entry in the solution matrix y to be the initial conditions for the next step. This allowed the virus to be injected at each time interval. For the initial run, these time intervals were $[0, 1.99]$, $[2, 3.99]$ & $[4, end]$ (days), where end represented the end time of the observed data.

As previously mentioned in section 4.1, the ODE solution is stored in 2 matrices upon completion of *ode45*. As they exit the function, they are stored in a larger matrix that contains all of the other outputs from y. These are organised by their cell population where each column is a new parameter set and the rows are the populations at the respective entry in t. The final size of these matrices is $end \times N$. The primary cell population of interest is susceptible cells, their data is interpolated to the times the data was observed from Kim et al. (2011).

The next portion was to generate weights from this new computed data based on their similarity to the

original data. This was done using the inbuilt MATLAB function *mvnpdf* which took in the data, a column from the susceptible cell matrix and the variance σ^2. These weights for each data set was stored in a new column matrix W where each entry was the weighting from a multi variable normal distribution for each parameter set. These weights were normalised using the sum of the weights, and 1000 of the highest weights were randomly sampled and their corresponding parameters were extracted. To visualise this result, a Kernel Density Plot was constructed for each parameter, including σ to see the distribution the highest weighting parameter sets took as well as their variance. This was done using the inbuilt MATLAB function *ksdensity*.

4.4   Previous Study

As previously mentioned, the data collected by experiments conducted by Kim et al. (2011) had the described ODE model fit to each mouse. The results from Kim et al. (2011) had parameters fit for an ODE model in a paper from Jenner et al. (2018). Figure 3 below displays two sets data and the fit models for each mouse.
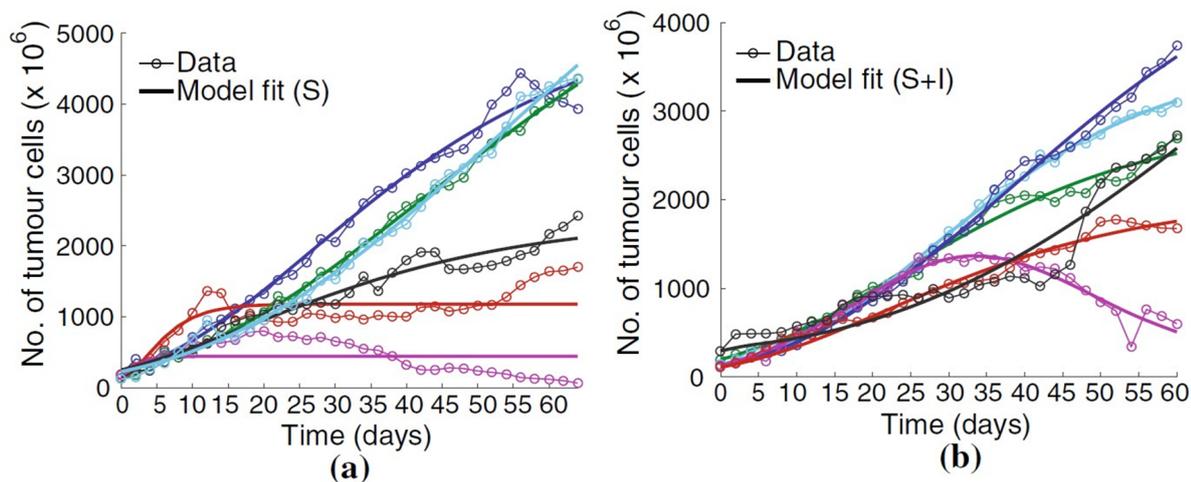
Figure 3: (a) Control Experimental Data (b) Naked Adenovirus Experimental Data

As the main component of the ODE model for controlled growth is a Gompertz growth term, we expect part a of figure 3 to display continuous growth as demonstrated by the green, navy and blue mouse. For the red and pink mouse, we see some abnormal growth patterns that don't follow the expected trends. We see a similar issue with the growth patterns for the red and pink mice when treatment is applied. For this reason, the red and pink mice were omitted from each set of Bayesian inference. Also to note, the mice across the two studies are **not** the same. They are colour coded in a similar way to identify each of the mice in the experiment. It isn't possible to monitor growth of a tumour in one mouse and add treatment to a newly added tumour elsewhere.

5 Results

5.1 Control Data Analysis

The following set of plots are Kernel Density Plots. These display the distribution of the highest weighted parameter sets for all 4 mice. On the x-axis is the values the parameters took and the y-axis represents the probability density.
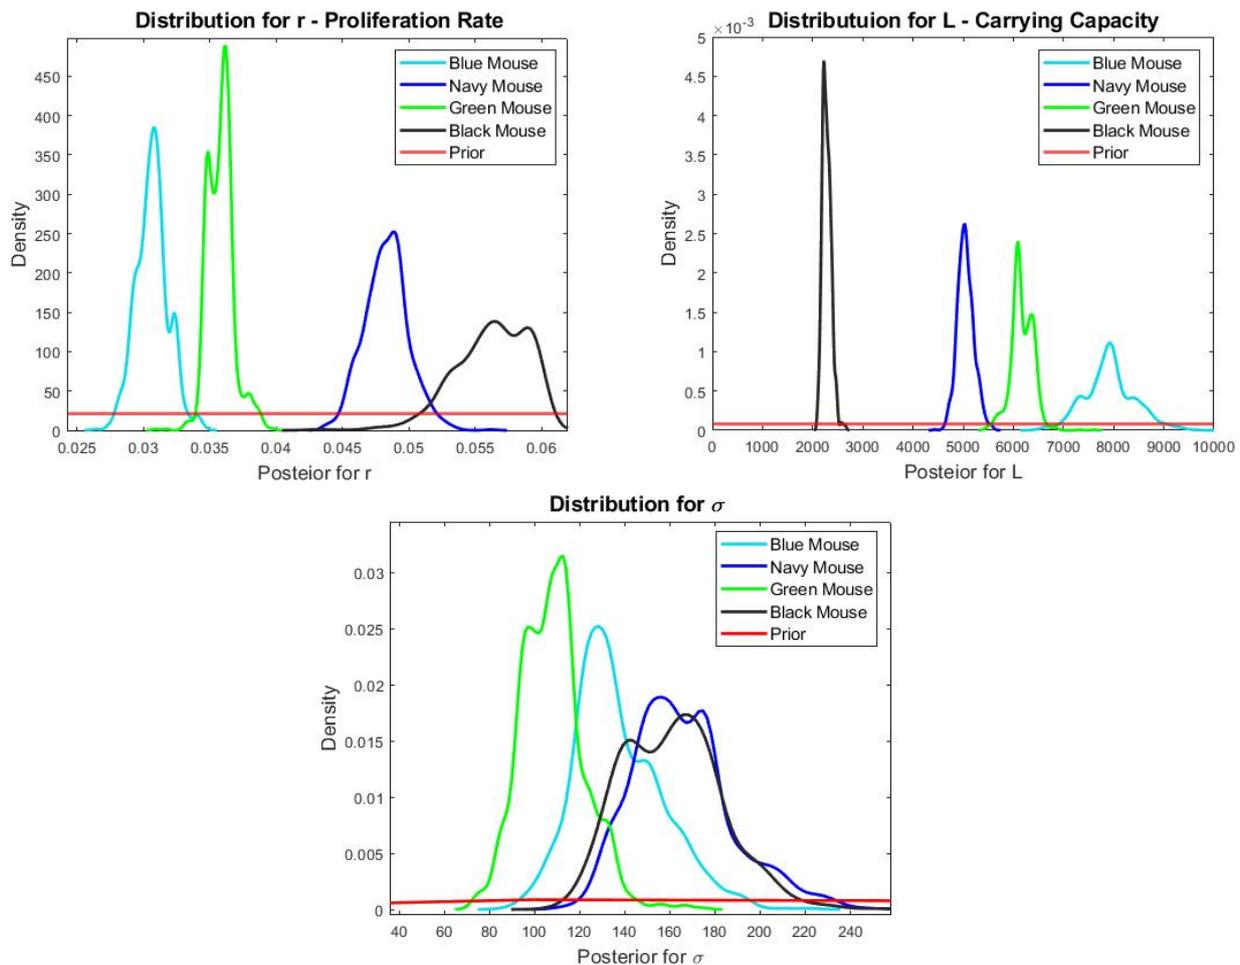

Figure 4: (a) Posteriors for proliferation rate, **r** (b) Posteriors for carrying capacity, **L** (c) Posteriors for variance, $\boldsymbol{\sigma}$

Figure 4 above displays the posteriors for r, L and σ. Part (a) displays posteriors for r. The red line displayed represents the uniform distribution of priors based on previous estimates. The smoother the distribution and the less range is has, the more certain we are in the distribution and the mean value. The black mouse does have 2 small peaks near its centre and a significantly longer lower tail compared to navy. The green and blue mice had distributions with small ranges, but there were several values that saw increases in density. The amount of smoothness observed in each distribution is most likely an artifact of not increasing the number of samples N to a high enough value due to computational constraints.

Part (b) of figure 4 displays the distribution for the posteriors for L. The red line again displays the prior distribution. In this figure, the black mouse displays the most certain result due to its smooth shape, small range but high density at a single point. This is closely followed by the navy mouse with a similar shape and size. The green and blue mice again had more bumps in their distributions but remained small in range. For Lm, they did also loose some density across the whole distribution.

Finally, part (c) displays the variance seen in each of the mice. The red line of the prior is the uniform distribution generated from a gamma function. Here the blue and green mice produced the most certain distributions. The black and navy mice did appear to have more variance in their distributions as evident by their width and low density. We are able to confirm this increase in variance for the blue and navy mice when looking at the comparison between data and the generated values.
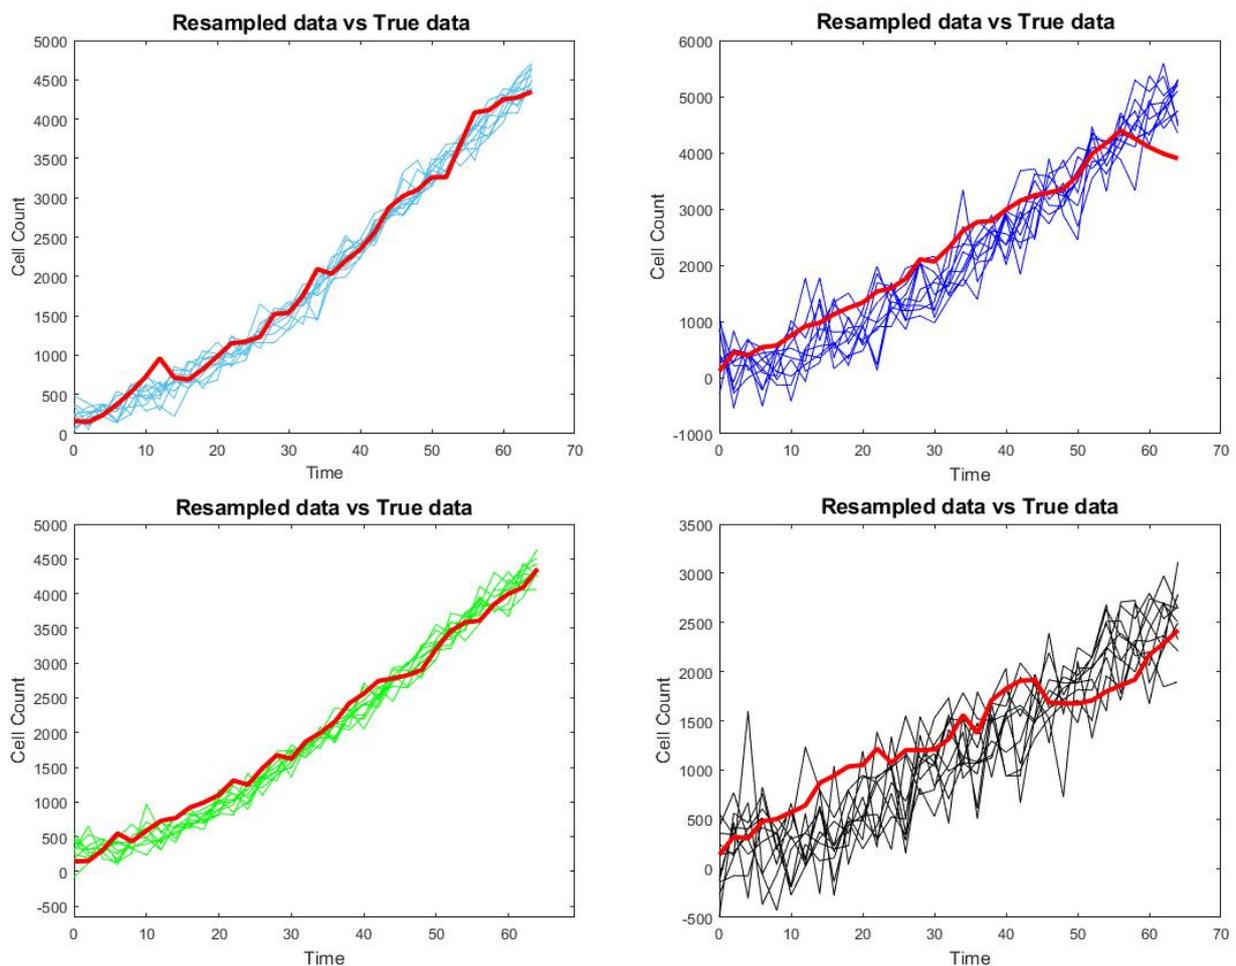

Figure 5: Comparison of data to 10 random parameter samples of (a) Blue Mouse, (b) Navy Mouse, (c) Green Mouse, (d) Black Mouse. The red line represents the original data set.

As seen above in 5, we can confirm the findings of figure 4c in terms of the variance. The blue and green mice have parameter sets that when simulated, produce data that follows the real data well. There are some

9

places in the true data, being the red line, that the parameter choices simply cannot replicate or are too noisy. We do confirm the increase in variance seen in the navy and black mouse here with a significant increase in the amount of noise present. The most prominent one being the comparison with the black mouse. This is a reasonable observation due to it having the least dense bur most sparse distribution in σ.

We can also compare the mean values in the posteriors captured in figure 4 with the original estimates for each parameter and mouse.

Table 2 - Comparison of mean posterior to estimated values

Mouse	r – Proliferation rate		L – Carrying Capacity	
	Previous Estimate	Posterior mean	Previous Estimate	Posterior mean
Blue	0.024	0.0308	10000	7923
Navy	0.043	0.0487	5300	5021
Green	0.03	0.0362	7300	6093
Black	0.05	0.0564	2400	2221

Table 2 displays the values estimated by Jenner et al. (2018) as compared to the mean posteriors for the control data. From this we can see the posteriors for blue resulted in a higher value of r, but a smaller value for L, these are 20% increase and decrease respectively. This is also seen in the navy mouse but to a lower degree, less than 15% change. Navy's comparison is much closer to the previous estimates in both parameters. Green has the most discrepancy in its value L at , but has a more refined values of r. Black also followed the higher r, and lower L trend but with the smallest difference of all the mice. Considering the Navy and black mice had the most variance but their values remained similar leads to a potential conclusion that the original estimates may not be the best suited to the data. On the other-hand, it could also mean we are generally unable to capture the true values of these mice to reproduce their data.

5.2 Adenovirus Data Analysis

A similar analysis can be performed on the data that utilises the Adenovirus treatment. The main differences in the model were the initial conditions now including viral particles and V_0 being non-zero. For the following figures, V_0 was defines as being 100 virus cells per injection unless specified. This also introduced a new parameter β, into the Bayesian Inference. This was another highly variable parameter across the mice in the fits by Jenner et al. (2018) and is the infectivity rate of the virus to susceptible tumour cells. As seen previously in figure 3, the red and pink mice had their data omitted due to abnormal growth patterns. The colours were randomly assigned to each mouse and the colours f mice to be discussed are not the same mice as section 5.1.
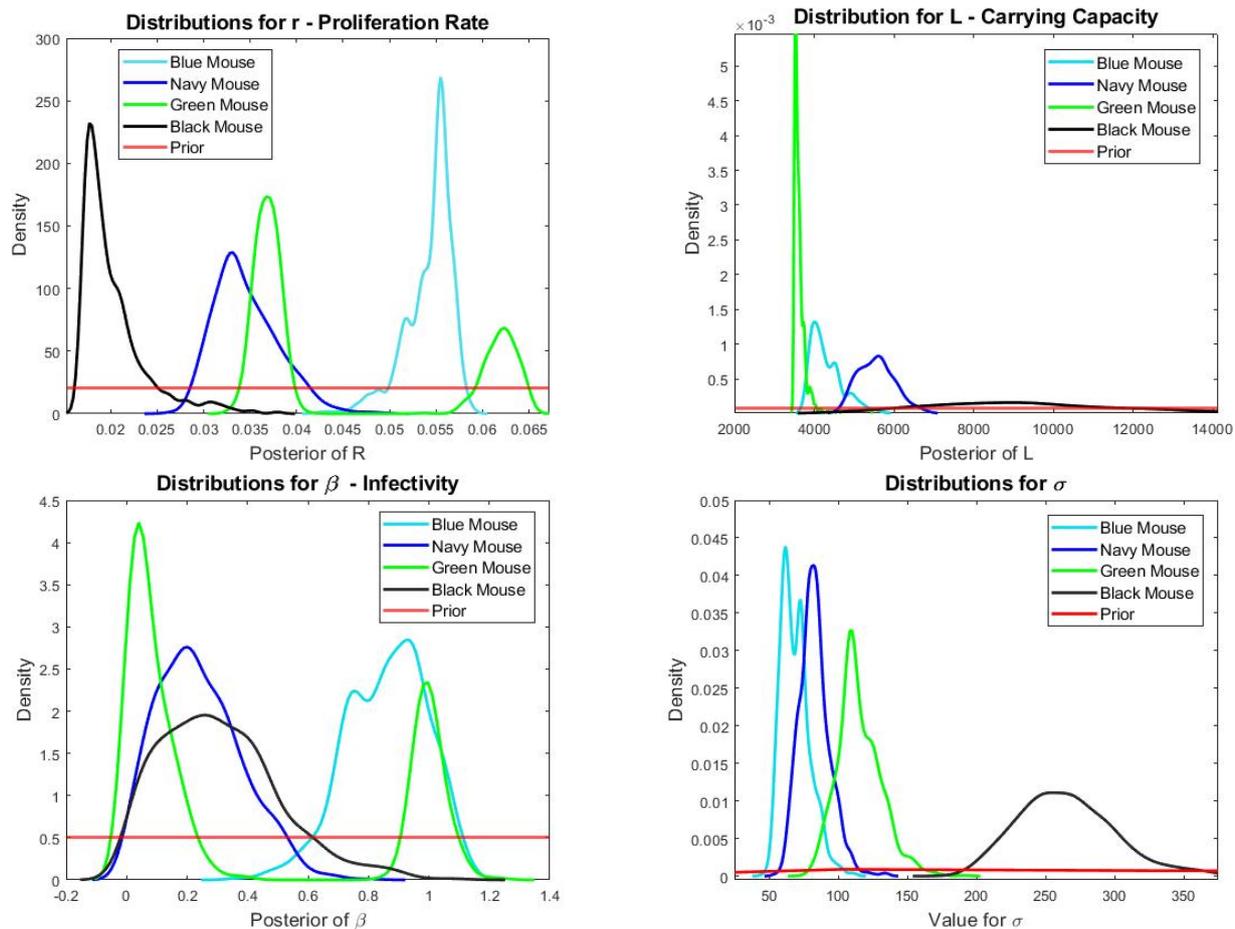
Figure 6: (a) Posteriors for proliferation rate, **r** (b) Posteriors for carrying capacity, **L** (c) Posteriors for Infectivity Rate, **β** (d) Posteriors for variance, **σ**

Above in figure 6, we can see the posteriors for each mouse and parameter, including variance σ. In part (a) of this figure is the posteriors for proliferation rate, r. While the black and navy mice produced relatively smooth distributions, black did have less range and was more dense indicating a more certain estimation. The blue mouse also had a dense peak but there were several bumps which suggests there were multiple values of r that were repeated. The green mouse is the most different as we can see two distinct peaks in the distribution. Individually they are both relatively dense, but their presence indicates it may be multi-modal.

We can confirm this potential multi-modal result for green when looking at part (c) of figure 6, the infectivity rate β. There is also two very distinct distributions forming. This suggests for green there are 2 sets of parameters that are able to achieve the same result in modelling the growth compared to the true data. The navy and black mice produced almost identical distributions for β, however, navy had a slightly higher density at its mean. The blue mouse was the most different with 2 small peaks and being centred over a higher value of β. It is interesting to observe that the two peaks in the distribution for green lie in the same ranges as other distributions.

Part (b) of figure 6 displays the distributions for carrying capacity L. Black has the widest and least dense distribution seen in any posterior. This indicates that our methods may not have fully captured the true range of values that best describe the black mouse's growth under treatment. The blue and navy mice had similar distributions in terms of shape and density. Finally, the green mouse was the most certain in almost an exact value for L due to its single, very narrow and dense distribution.

The shape and sizes of the distributions for each mouse were accurately reflected in the distributions for variance σ. We can see that green, blue and navy have similar distributions and all had reasonable posteriors for the parameters The one for green solidifies the idea that it is multi modal. If this were not the case, its variance distribution may be wider and more flat like black is. Potentially due to its high level of uncertainty in its distribution for L, the amount of variance seen in the black mouse is increased. This is due to it being centred around a higher value of sigma, but also being wider and less dense compared to the other mice.

In all cases, N was increased from 3 million to 30 million to acheive this level of smoothness. Further increases would have been possible if not for a limit on computation time and power available.

We can summarise the findings from figure 6 below as Table 3, and compare the mean of the posteriors distributions to the values previously estimated.

Table 6 - Comparison of mean posterior to estimated values

	r – Proliferation Rate		L – Carrying Capacity		β – Virus Infectivity	
Mouse	Previous estimate	Posterior Mean	Previous estimate	Posterior Mean	Previous estimate	Posterior mean
Blue	0.057	0.0554	3700	3992	1.5	0.926
Navy	0.036	0.0328	5600	5643	0.2	0.200
Green	0.05	0.036/ 0.0627	3000	3538	0.85	0.040/ 0.996
Black	0.022	0.0177	10,000	8871	1.2	0.258

Above in Table 3, we can see the comparison in estimates for each parameter and mouse. Overall, the blue mouse produces similar estimates with its most significant difference being a decrease in the blue β. Navy had an exact match for β and was within less than $\pm10\%$ for all estimates. Green was more complex due to its multi-modal distributions, however in r, if the mean of the two estimates are taken we achieve $r = 0.04935$, which is a 1.3% error. The same small error cannot be said for the values in the β posterior with close to 40% error. The value observed in L is also within $\pm20\%$. Black mouse had the most significant change being its estimation for β with a close to $5\times$ decrease in the value. Its other estimates were also within 20% of the previous estimates. There also was a general trend in the change in relationship between r and L. We know from previous investigation in this project that these two variables form a near perfect polynomial relationship. This would support the reason why a decrease in r, (which all mice saw) resultsin an increase in the value of L.

5.3 Treatment Analysis

With these posterior values now generated, it is possible to simulate a potential patient response to varying treatments. Based on the Adenovirus treatment regime, 3 alternatives were developed. This included increasing the number of viral particles per dose from 100 to 200. Another included increasing this further to a total of 500 cells per injection. The final was to maintain the same 100 cells per injection, but instead of 3 doses every 2 days, this was altered to be 1 dose per day for 7 days.

To simulate these different treatments, the re-sampled parameters that formed the posteriors were grouped into random parameter sets and run though the ODE. The number of susceptible cells at each time was stored in a larger matrix and these values were averaged at each time point. This was repeated for each new treatment regime by altering the value of V_0 or the number of time increments the model is run with to allow more injections.
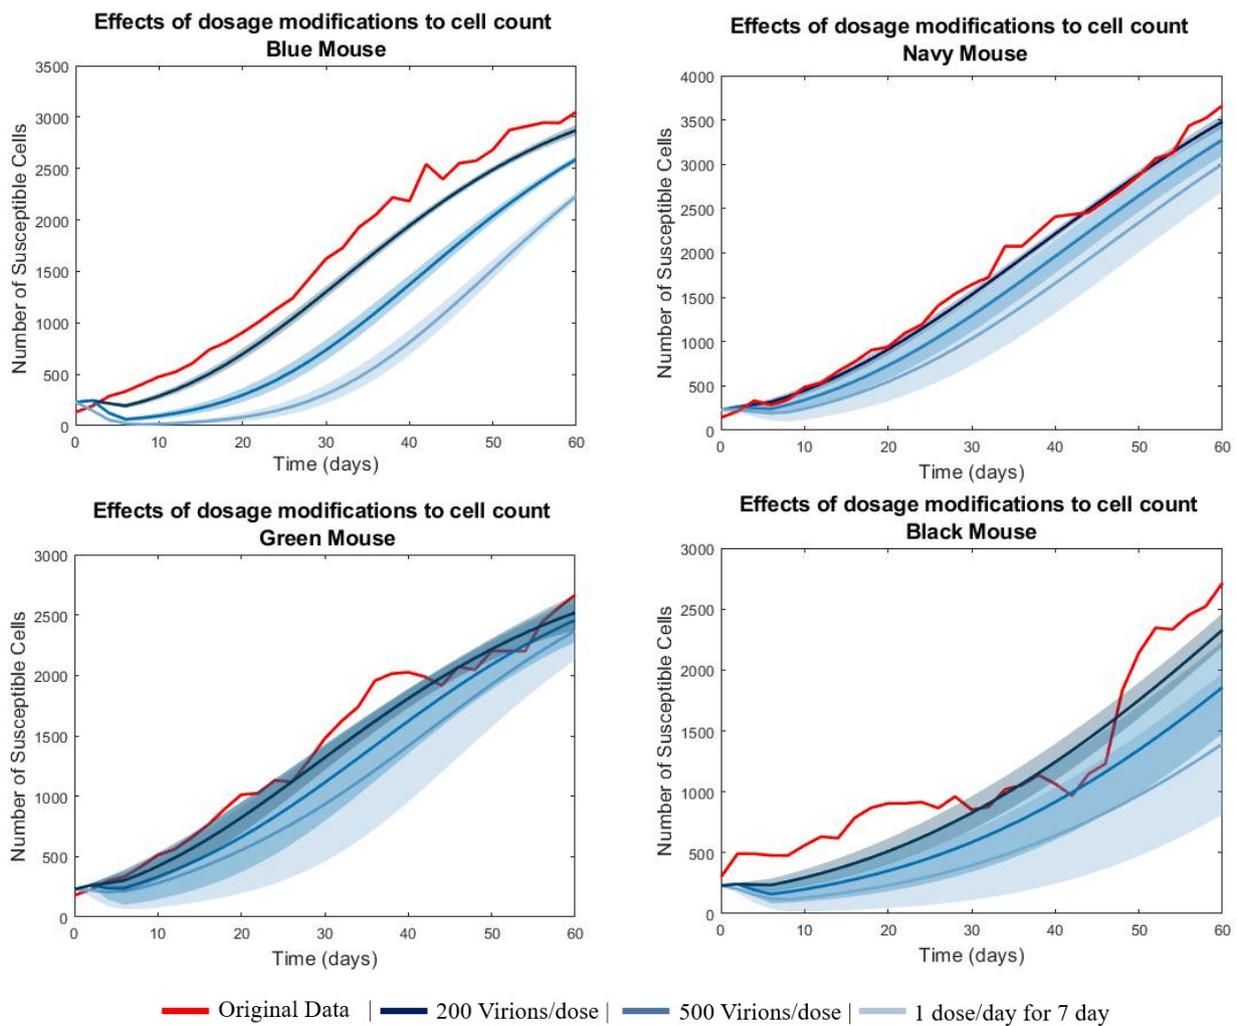

Figure 7: Number of susceptible cells over time under varying treatment regimes

Above in figure 7, we have the results of altering the treatment across all 4 mice as well as the variance associated with each change. Each plot displays the amount of time in days on the x axis, and the number of susceptible tumour cells on the y axis. The red line indicates the data from the current treatment regime. In all four mice, increasing the number of doses received in total from 3 to 7 and increasing the frequency of doses made the most significant impact on the number of susceptible cells. Increasing the number of virions in each dose did also decrease the number of susceptible cells over time, but not to the extent of more doses.

It can also be observed the amount of variance in each of these cases. The blue mouse had the least amount of variance with these dosage changes which would make it a hyper-ideal patient. The navy mouse was second in having reduced variance, it also did not see as significant of an impact on tumour size when the treatments were altered. The green mouse experienced more variance, but similarly to the navy mouse, a smaller impact when dosages were changed. The black mouse however saw similar effects as the blue mouse but with significantly more variance. It does have the most abnormal growth pattern already, and this may have had an effect on variance.

6 Future Work

If this information were to be utilised in the progression of this technology in humans, you would be more inclined to go with the blue mouse initially. The group of mice utilised for this experimental data from Kim et al. (2011) were specifically bred to genetically identical mice. As seen previously, there is still a significant amount of variance in their response to treatment as well as general growth patterns. If this is the level of variance experienced here, imagine the increase in variance seen in a population of humans. Already in other medical treatments there is a lot of variation in response to particular medications and other therapies, and it is increasingly difficult to capture the entire population. In this case, we would be inclined to model responses based off the green and black mouse due to their increase in variance.

Future work in this area is very possible. It would be first ideal to investigate if it is the number of doses received or the frequency of treatment that has the greatest impact on the treatment response. This would allow us to further model different treatment regimes and see if there is an ideal one for this cohort of mice. We also have access to tumour growth data in mice for lung and ovarian cancer. It would be interesting to see if there is a difference in the mean posteriors for each mouse and parameter. This would allow us to build more of a database of potential responses for a mouse cohort.

AMSI

7   Conclusion

To conclude, based on prior experimental data and ODE model, we are able to utilise Bayesian Inference to generate parameters for each mouse based on their individual growth data. Each of the means of the posteriors were comparable to the previous estimates for these parameters. This provides further support to these estimates. In using these new parameters, it was possible to simulate a range of new treatment options in efforts to predict the mouse's treatment outcome. In all cases, more frequent smaller doses of a longer period of time proved to have the most significant reduction in the number of susceptible cells. Further research in this area includes investigating the true relationship between frequency of doses and duration of treatment. This would allow new tumour data sets that are available for liver and ovarian cancer to have this process applied to and results compared. The ultimate goal of further works would be to extend this technology from mice to a human cohort and specialise parameters based on patient characteristics such as age or medical history. This is far into the future but is a powerful piece of work to be undertaken for the improvement of cancer treatments.

8   References

Hrouda-Rasmussen S (2021) "What is Bayesian Inference?" Towards Data Science
`https://towardsdatascience.com/what-is-bayesian-inference-4eda9f9e20a6`

Kim PH, Sohn JH, Choi JW, Jung Y, Kim SW, Haam S, Yun CO (2011) "Active targeting and safety profile of PEG-modified adenovirus conjugated with herceptin". Biomaterials 32(9):2314–2326

Jenner AL, Yun CO, Kim PS, Coster ACF (2018) "Mathematical Modelling of the Interaction Between Cancer Cells and an Oncolytic Virus: Insights into the Effects of Treatment Protocols" Society for Mathematical Biology 80:1615–1629

Pitone ML (2022) "Adenovirus". The Nemours Foundation - Nemours Children's Health.
`https://kidshealth.org/en/parents/adenovirus.html`

Tjørve KMC, Tjørve E. (2017) "The use of Gompertz models in growth analyses, and new Gompertz-model approach: An addition to the Unified-Richards family." doi: 10.1371/journal.pone.0178691

Trager, Megan H., Larisa J. Geskin, and Yvonne M. Saenger. (2020) "Oncolytic viruses for the treatment of metastatic melanoma." Current treatment options in oncology 21 : 1-16.

Wares JR, Crivelli JJ, Yun CO, Choi IK, Gevertz JL, Kim PS (2015) "Treatment strategies for combining immunostimulatory oncolytic virus therapeutics with dendritic cell injections." Math Biosci EngMBE