

Exploring the Change of Malarial Resistance in Response to Repeat Antimalarial Treatments Through Agent Based Computer Simulations

A.Walker

A.Precious

R.O'Donnell

L.Baran

*School of Computing
Teesside University
Q5045715@tees.ac.uk*

*School of Computing
Teesside University
Q5068888@tees.ac.uk*

*School of Computing
Teesside University
Q5273477@tees.ac.uk*

*School of Computing
Teesside University
Q5127950@tees.ac.uk*

Keywords: NetLogo, Malaria, Antimalarial Treatment, Artificial Intelligence, Agent Based Model

Abstract

Following work done by a paper exploring the evolution and impact of antimalarial drug resistance, two computer simulations were developed using NetLogo. Both models proved the hypotheses set out from the original paper (Christian Nsanzabana, et al, 2010). The systems demonstrated emergent properties that confirmed the results put forth by the source paper. Model findings support the qualitative results of the source paper in suggesting that a link exists between the application of antimalarial treatments and the increasing resistance of malaria.

1. Background

1.1. Source Paper

The work carried out by the team in this paper is based on the source paper “Quantifying the Evolution and Impact of Antimalarial Drug Resistance: Drug Use, Spread of Resistance, and Drug Failure over a 12-Year Period in Papua New Guinea” (Christian Nsanzabana, et al, 2010).

In addition to the paper put forth by Christian Nsanzabana, et al (2010), the work displayed in this report was also informed by papers from Jan Engelstädter (2017), Anabela Simões , Ernesto Costa (2000), Klein E.Y (2013) and White N.J (2004).

The source paper was produced based on work conducted as part of the “Malaria Vaccine Epidemiology and Evolution project”, which studied residents within the catchment area of the Kunjingimi health centre in the East Sepik Province of Papua New Guinea between the years of 1991 and 2002.

Within this province, the Kunjingimi health centre was the sole provider of healthcare, effectively creating a closed population ideal for study. During the study period, data pertaining to the diagnosis and treatment of malaria from patients visiting the health centre was collected.

From the clinical records, molecular analysis of blood samples, and surveillance carried out by the research team, the paper identified three key components that lead to failure of the antimalarial drug:

1. “A high level of antimalarial drug use within the community, creating the driving force for resistant mutations.”

2. “The spread and propagation of resistant mutations within a population due to evolutionary pressure.”
3. “Failure of the Antimalarial drugs to treat subjects due to the increased resistance possessed by Malaria against Antimalarial drugs.”

The paper investigated the relationship between the three factors in order to quantify the impact of drug usage in driving mutations within the case study area of Papua New Guinea. This paper focuses on modelling evolutionary resistance growth through allele and haplotype frequencies rather than analysing patient bloods to discover the mutation prevalence. The study revealed that treatment failure rates within the case study multiplied by 3.5 times between the years 1996 to 2000, but then decreased after a drug policy change in the area.

Though the paper does not explicitly state it as an outcome of the investigation; it suggests a link exists between the the Kunjingimi health centre staff’s liberal application of antimalarial drugs to any patient that exhibits fever symptoms, and the strength of the resistance of the malaria against antimalarial treatments. The paper also states that despite the introduction of drug application policy changes and the introduction of a revised drug being capable of drastically reducing treatment failure rates, they still observed a rapid increase in the number of resistance markers and clinical failures.

1.2. This Paper

Based on the work by Christian Nsanjabana, et al, (2010), the team decided to investigate the effects of repeated treatment cycles in which new strains of the antimalarial treatment were administered to an populus infected with malaria, similar to that witnessed in the Papua New Guinea case study. To gather the data for this investigation, the team decided to build an agent based computer model using the NetLogo environment.

For the malaria model, the team worked to test the following hypotheses: “The effectiveness of antimalarial treatment against malaria will decrease with time as the malaria becomes resistant”, and “Each subsequent application of a new strain of antimalarial treatment will have a lower overall effectiveness than prior applications due to increased malaria resistance”.

For the human model, the team worked to test the hypothesis that: “The frequency of application of new strains of antimalarial treatments has little long term effect on the population of malaria within an ecosystem”.

2. Method

Two patch based NetLogo models were produced, both models utilise the same back end for managing the genetic mutation and reproduction algorithms, the antimalarial treatment success calculations, and the logging subsystem. One model (referred to as the malaria model) looks specifically at the malaria cells independently of other factors with each patch representing a malaria cell. The other model (referred to as the Human model) looks at how the malaria interacts with a human population with each patch representing a human.

Both models implement a binary genome for their malaria cells and treatment. Each binary genome is a randomly generated list of 32 bits. Lists were chosen to store the genome as in built NetLogo functions allowed for easy comparison of different genomes, as can be seen by the following code snippet utilising the *map* function.

```
ask patches with [ state = "alive" ]
[ set comparator (map = genome treatment-genome)
  if (( length filter [ i -> i = true ] (comparator) /
    length comparator)*100 )<treatment-effectiveness
  [ set state "dead" ]
]
```

Figure 1 - The core code responsible for treating malaria

The treatment of malaria is the same across both models and is considered effective when the genome of the treatment is of a certain similarity to the genome of a cell. This similarity is set by a slider, that can be adjusted at the GUI level of the model, named “treatment-effectiveness”. The similarity is arrived at by comparing each bit in the list of the malaria genome to the corresponding bit in the treatment genome and noting if they match as shown in the diagram below.

```
Malaria : 1 0 0 1 1 0 1 0 0 1 1 0 0 0 1 1 1 1 0 0 1 0 0 1 0 1 0 1 0 1 1 1
Treatment: 1 0 1 1 1 0 1 1 1 1 1 0 1 0 1 0 0 1 0 1 0 1 0 0 1 0 1 1 1 0 1 1
          T T F T T T T F F T T T F T T F F T T F F F T F F F F T F F T T
```

Figure 2 - Malaria Treatment Example

Figure 2 displays the 32 bit binary genome of a malaria cell and the 32 bit binary genome of a treatment. The 3rd row of the diagram indicates whether or not the malaria genome bit at that index matches the corresponding bit from treatment genome, if the two values match, the 3rd row gets a T for true, if they do not match the 3rd row gets an F for false. The number of Ts are then summed and compared to the value held by the “treatment-effectiveness” slider. If the total number of Ts is higher than the value held by the “treatment-effectiveness” slider then the cell is considered to be treated.

2.1. Modeling Malaria Cell Genetic Mutation to Antimalarial Drugs

The malaria focused model looks specifically at how the malaria cells behave when the human element of the model is removed. The malaria model can be thought of as a petri dish with each patch representing a malaria cell as can be seen in figure 3.

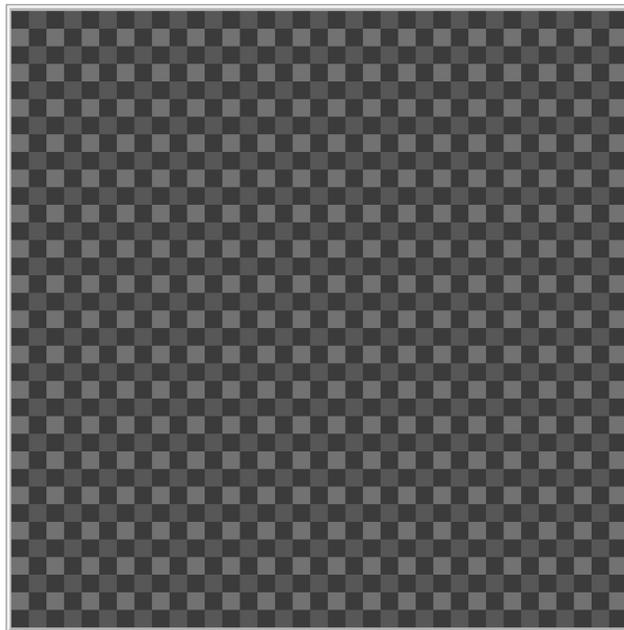


Figure 3 - NetLogo model in completed state

Upon setup, each patch is given a randomly generated 32 bit genome and a random treatment is generated. A new random treatment can be introduced at a certain number of ticks to further combat the malaria. The model has completed running when the malaria cell count has reached the original population count. The model implements a state machine pattern based on the following state diagram:

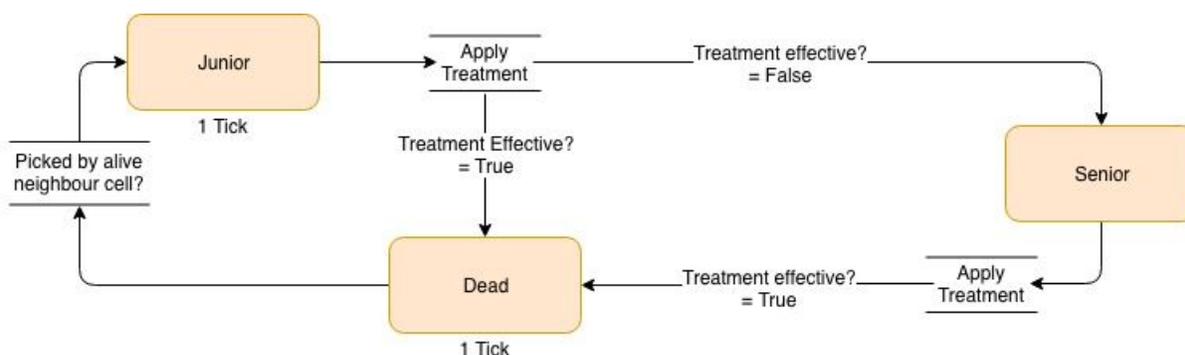


Figure 4 - Malaria State Machine Model

Looking at figure 4, it can be seen that the start of a malaria cell life cycles begins in the “Junior” state. The primary characteristics of a junior malaria cell are that it is infertile, meaning it cannot reproduce until it has matured and transitioned into the “Senior” state. This life cycle model was based from information put forward in a paper by Klein, E.Y (2013). An additional feature of this system allowed the amount of malaria cells produced in a single tick to be counted, as without this state transition the new malaria cells would reproduce in the same tick in which they were conceived, making it difficult to track reproduction per tick.

Once a junior cell has been created a treatment is then applied to see if the cell will survive. The treatment is not only applied to the cells in the Junior state it is applied to all alive malaria cells, as can be seen in the state transition diagram following the Senior cell. The reasoning for this is that cells in the Senior state may also be treated if the antimalarial treatment has been changed since it transitioned from the Junior state. If the treatment is effective then the cell transitions into the Dead state as can be seen by the state transition model.

Once a cell has reached the Senior state it has gained the potential for reproduction. The reproductive processes occurs when the cell is chosen to reproduce on the current tick. The way in which a cell is chosen is tied to a slider on the GUI level of the model allowing the user to select a percentage value of the alive cells to reproduce. If the cell has been randomly selected to reproduce, it will select one of it's 8 neighbors, as shown in the diagram below.

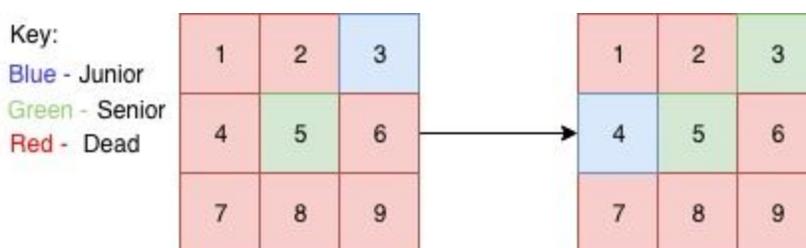


Figure 5 - Neighbor Selection

In the scenario provided in figure 5, patch 5 represents an alive malaria cell in the Senior state that has been chosen to reproduce. The first thing it must do is select a neighboring patch to reproduce to. The neighboring patch must be in the Dead state in order to be selected for a new malaria cell to take its place, this means that the selection of patches that can be picked are: patch 1, patch 2, patch 4, patch 6, patch 7, patch 8 and patch 9. Patch 3 cannot be selected as it is currently occupied by an alive malaria cell in the Junior state. Once a dead neighboring patch has been randomly selected, in this case patch 4 has been selected, the cell that is reproducing (patch 5) runs its own binary genome through a mutation function that outputs a new mutated genome that will be given to the randomly selected neighbor patch (patch 4). Patch 4 will then transition into the Junior state and be assigned its new binary genome as can be seen by the diagram. The diagram also demonstrates the transition of patch 3 starting at the Junior state on the left and transitioning into the Senior state on the right. This process takes place over 1 tick.

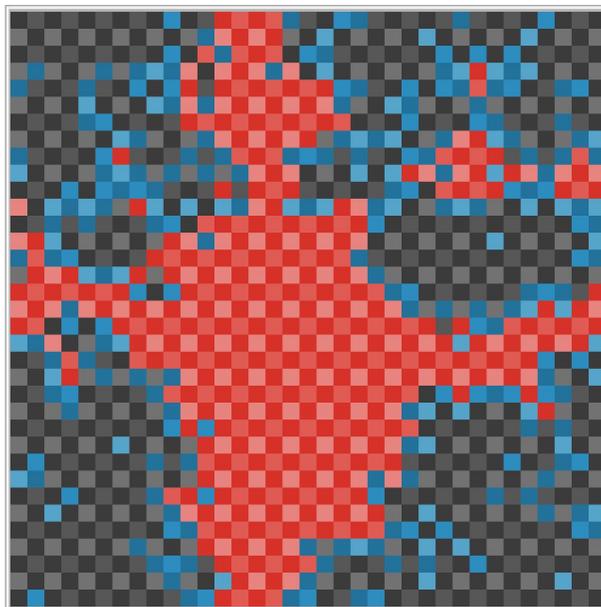


Figure 6 - Grouping of Malaria cells from a NetLogo simulation

Figure 6 demonstrates a property of the model in which the malaria cultures group together and grow. This behavior is to be expected when taking into account the breeding methodology in which a senior malaria cell can only reproduce to an empty space within the 8 neighboring cells.

2.2. Modeling Malaria Spread Between Humans

The human model was built to resemble malaria infection spread, its treatment and its behavioural changes among the human population. The asexual reproduction and mutation methods used are identical to the previous model. Each patch in the NetLogo environment represents a human individual. Each human has values representing age and its health state. Health state is represented by an appropriate colour as can be seen by the following state machine diagram.

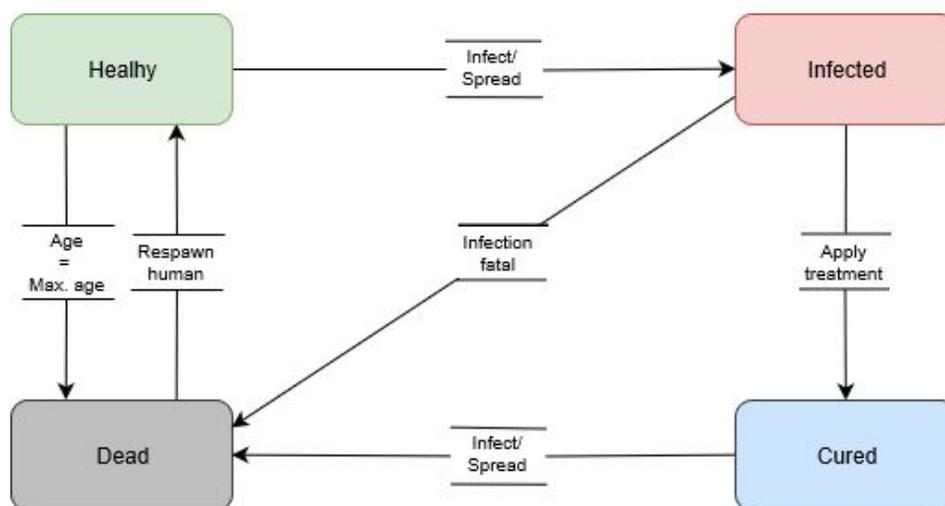


Figure 7. Human state machine

Initially all patches start in healthy state - not being infected or treated from malaria. Individuals can die of natural causes after reaching a certain age represented by ticks. Individuals are infected with malaria at random by the “Infect” button on the NetLogo GUI. After more individuals are infected malaria will start spreading through the population. Every infected individual is represented by a red colour and has an infection length meter. If the infection is not treated within a specified amount of time the individual dies, as evident in the source paper (Christian Nsanzabana, et al, 2010).

An antimalarial treatment is applied per tick to the whole population. After a specified number of ticks, a new treatment is introduced to the system in order to combat the malaria strains that have become immune. The number of individuals it is possible to treat depends on the user defined treatment success rate, available to the user on the GUI level of the NetLogo model. The treatment pattern used is identical to the malaria cell model discussed above.

After an individual dies, whether the cause is malaria or age, it is visually represented by marking the corresponding patch in black. After a specified amount of time a new healthy individual takes its place. The time taken before spawning a new human is modifiable but for keeping a consistent population model it is advised to keep the time at 1 tick.

2.3. Modelling Breeding

To accurately model asexual reproduction of malaria, research was conducted into the process by which a genome is mutated during reproduction. From this investigation, the team discovered a paper detailing the genetic algorithms used for transposition (Anabela Simões , Ernesto Costa 2000) in asexual reproduction in a binary genome. It was this mutation model the team decided to implement in their NetLogo model.

However, due to strict time constraints and lack of experience and limitations of handling and manipulating large datasets with NetLogo, the team opted to design a simplified transposition algorithm whereby the the genome was simply reversed rather than fully pattern shifted.

```

to-report simple-swap-set-length-breed [genomeIn setlength]
  let glength length genomeIn - 1

  let swap-start-point random(glength)
  let swap-end-point 0
  ifelse((swap-start-point <= (glength - setlength))) [
    set swap-end-point (swap-start-point + setlength)
  ] [ set swap-end-point glength ]

  let swap-partion reverse (sublist genomeIn swap-start-point swap-end-point)

  let swap-index 0
  let local-index swap-start-point
  let comp genomeIn

  loop [
    set comp replace-item local-index comp (item swap-index swap-partion)

    set swap-index swap-index + 1
    set local-index local-index + 1

    if local-index = swap-end-point [
      if length genomeIn != length comp [show "HAS BROKEN"]
      report comp]
  ]
end

```

Figure 8 - The final NetLogo procedure used to perform a mutation on a genome

The procedure shown in figure 8 starts by selecting a random position within the binary genome as a starting point, then calculates an endpoint using the “setLength” end bit. The procedure is designed to prevent out of bounds errors by calculating whether the endpoint is outside the genome length, and if so limiting the endpoint to the maximum length of the genome. The procedure then reverses the bit order of the selected partition to mutate the genome. This part of the procedure takes place within the loop statement.

2.4. Logging Model Results

During development of the NetLogo models, the team decided it would be beneficial to have greater diagnostic utilities to debug model behaviour. To that extent, a model agnostic logging library was written to an .nls file that could be imported into the teams NetLogo projects. This allowed the same logging behaviour to be used both in the human and malaria models. Below is the main procedure of the logging library:

```

to !logg-fwrite [#tick #name #data]
  let body fput #tick #data
  file-open (word #name "_" logg-date ".csv")
  file-print csv:to-row body
  file-close
end

```

Figure 9 - The core code of the logging library written for the human and malaria models

The above code snippet in figure 9 shows the function “!logg-fwrite”. Note how the function accepts parameters for the current tick number of the execution and the name of the file to log to as well as the dataset. The procedure also consumes an internally generated date-time, named in the code as “logg-date” which is appended to the file name.

Logging the tick number as part of the dataset printed to the file allowed the team to mitigate issues arising from where the same file was written to multiple times per tick, such as when the logging function was called from within an agent or patch each time certain interactions between agents took place.

By accepting a name parameter, the logging library allows multiple log files to be written to for each run, holding different diagnostic recording data from the execution and making the recorded data easier to traverse and understand.

Later on in the projects cycle, after the models codebase was completed, the team made use of BehaviourSearch: a component of the NetLogo ecosystem that allows many instances of the a model to be ran in unison. The executable supports functionality to seed each model instance with different parameters, allowing for a range of inputs to be tested. This feature was used extensively by the team to generate the statistical data used to produce the graphs used in the Results section of the report.

3. Results

Using the NetLogo Behavior Space functionality, the group was able to repeat the experiment multiple times using different input variables. This section of the paper looks at the results of both models, linking back to the hypothesis of the original paper.

3.1. Results From the Malaria Resistance Model

Given certain variables for the malaria model, some interesting emergent properties can be seen, particularly when looking at the “treatment-effectiveness” input variable. The graphs shown below demonstrate the journey of results when the treatment-effectiveness slider is changed, starting at 50% and ending on 76%, (the percentage at which all malaria cells are instantly treated) with an increment of 2% each run. In the context of the graph, a blue line represents the count of malaria cells in the Junior state, the red line indicates the amount of malaria cells in the senior state and the grey line represents the number of dead malaria cells. The X-axis represents the time passed in ticks and the Y-axis represents the count of malaria cells.

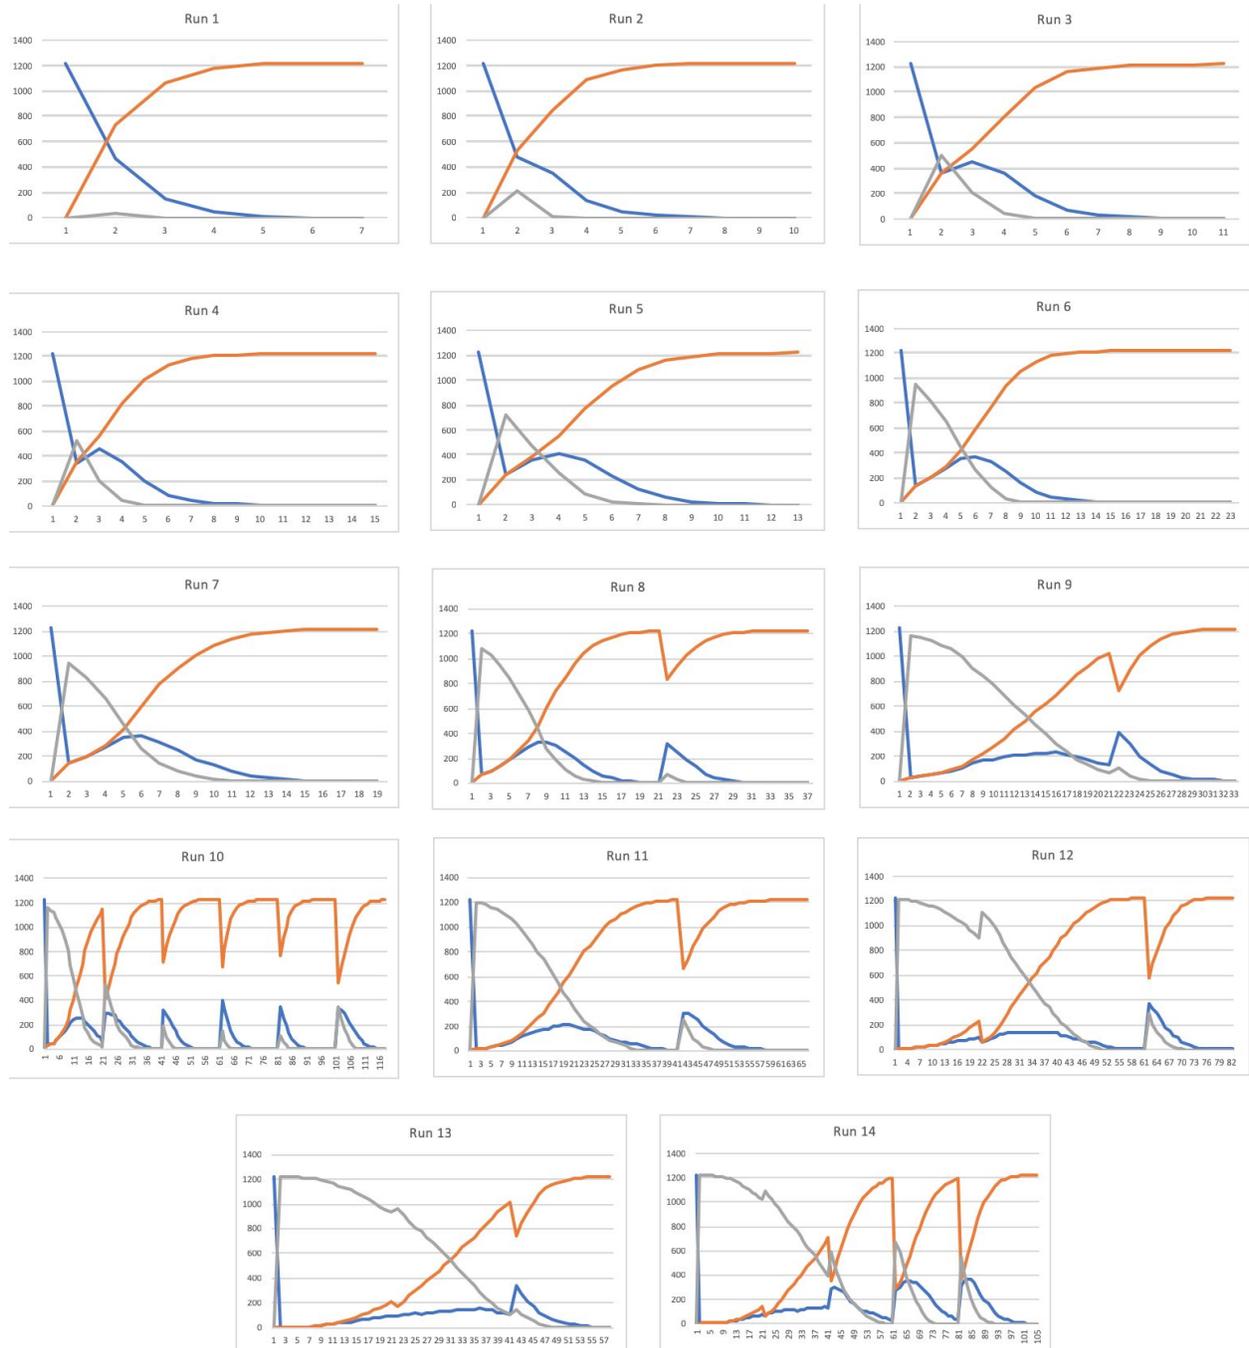


Figure 10 - Treatment Effectiveness Increasing per Run

As the Treatment Effectiveness value increases, the grey peak showing how many malaria cells that have been treated grows. This peak continues to grow rapidly throughout the graphs until run number 8 where a second peak can be seen. This second peak marks the introduction of a new treatment which has been introduced every 20 ticks. It is at this point where the first emergent property can be seen; the peak for each subsequent treatment never exceeds the height of the previous peak (with exception of Run 10, discussed later in the document). This implies that with each new treatment introduced, the treatment

becomes less effective at killing the malaria cells until the treatment has little to no effect at all, thus the malaria has become resistant to treatment.

Run number 10 looked particularly interesting as the malaria appeared to reach an equilibrium where the malaria population recovered by around the same amount that was killed each new treatment cycle. This system remains stable until the tipping point at the 100 tick mark where the last treatment was introduced. Even though this treatment killed off more malaria cells than the previous 3 treatments, (this can be seen by the larger than normal dip of the red line indicating the malaria population in the senior state), the malaria had mutated to a point where it could recover and regain the population of malaria cells to 100% before the next treatment is introduced. This is interesting because it demonstrates the emergent property that the malaria cell population has gained, where it has become increasingly resistant to the treatment with each subsequent treatment applied to the system.

The original paper hypothesised that: “The effectiveness of antimalarial treatment against malaria will decrease with time as the malaria becomes resistant”. According to the results gathered from the malaria based model, it can be concluded that this is indeed the case, indicated by the following graph:

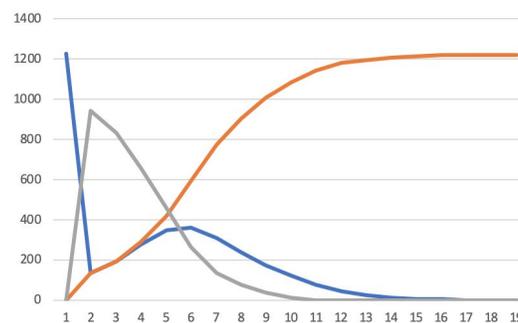


Figure 11 - Treatment Effectiveness Decreasing Over Time

As with the previous graphs, the blue line represents the count Junior cells, the red line indicates the count of senior cells and the grey line represents the number of dead malaria cells. The sharp spike in the grey line at tick number 2 indicates that the treatment has been introduced to the system and around 875 malaria cells have been affected by the treatment. The effectiveness of the treatment begins to decrease over time as is indicated by the fall of the grey line and the increase of the red line, indicating malaria cells in the senior state. Looking at the previous results generated by the malaria model it can be seen that they all follow the same pattern, indicating that the hypothesis of the source paper to be true.

The second hypothesis set out to be tested by the malaria model was that: “Each subsequent application of a new strain of antimalarial treatment will have a lower overall effectiveness than prior applications due to increased malaria resistance”. When looking at the results generated by the malaria model, it can be concluded that this is indeed the case.

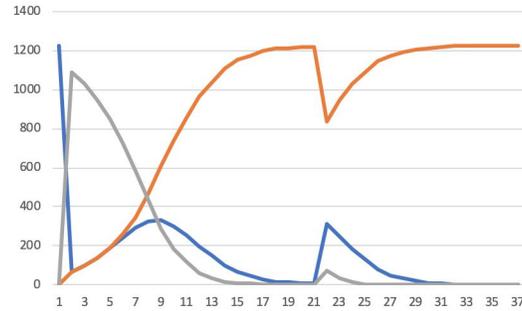


Figure 12 - Treatment Effectiveness Decreasing with Substiquent Treatments

As indicated by the above graph, the first treatment is introduced at tick number 1 and is effective against around 1050 malaria cells in the system. The second treatment is introduced at 20 ticks at the effects can be seen with the spike at tick number 21-22. Although this treatment is effective against a number of malaria cells, the overall effectiveness is far less that that of the original treatment. This effect can be seen in the previous results where the subsequent treatment peaks generally do not exceed that of the previous peak. Due to the fact that the new treatment being introduced has an element of randomness, it can occasionally produce a significantly effective treatment that may match the previous treatment in effectiveness, causing some outliers for this rule. However the vast majority of tests follow this rule, indicating that the hypothesis of the original paper is correct.

3.2. Results From the Human Malaria Transmission Model

For each run of the human model approximately 8 to 20 randomly selected individuals were chosen to be infected. In most cases similar results were observed. The malaria spread factor could either be set too high or too low. If it was too high (6+) malaria quickly spread through the whole population leaving little chance for the treatment to produce a significant effect. With the malaria spread factor set too low, quick eradication of the the malaria was observed due to the high treatment effectiveness and repressed ability for the malaria to mutate. These results did not seem to resemble the real world population as described in the case study.

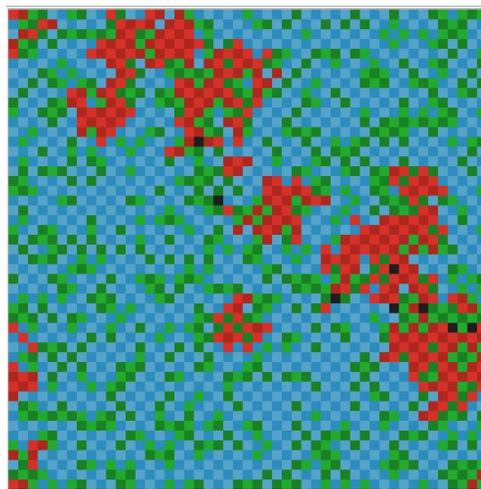


Figure 13 - Infection isolation

Based upon the team's findings, the infection-spread and treatment-effectiveness factors were most physically accurate when capped around 6.3 to 6.7 and 5.1 to 5.5 respectively. Emergent properties were exhibited in test runs with a treatment strength higher than the value of 5.5, whereby the treated population was isolated from the infected communities as represented by figure 13. This caused the malaria infection to be eradicated from the populus.

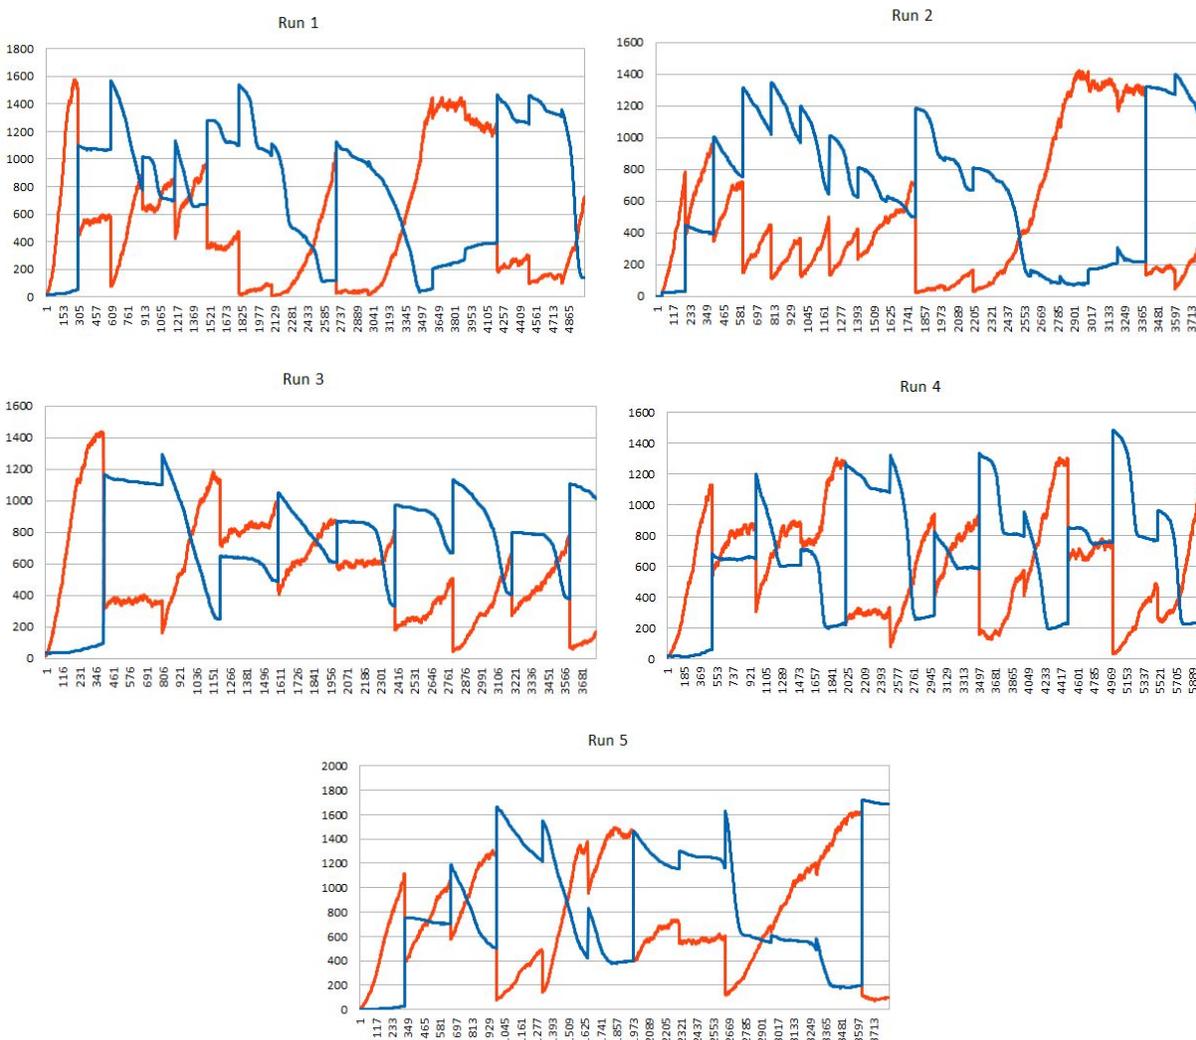


Figure 14 - Human Testing Results

The behavior represented by the graphs in figure 14 is similar to that of the ones shown in the malaria model, it can be seen in figure 14 that the lines show a repeating increase in infected human population, despite the introduction of new treatment cycles. However, the line profile of the graphs in figure 14 is irregular and largely sporadic, suggesting the quality of results attained is questionable. The human model would largely benefit from implementing the 4-attribute based breeding. Due to the interplay of factors within the model, it is difficult to confirm the credibility of the results.

3.3. Analysis of results

The general effectiveness of the models produced for testing the hypotheses under investigation was considered reasonably strong by the team. The malaria model generally exhibited the expected results, with the graphs showing a general downward trend in the effectiveness of a single application of antimalarial treatments at killing malaria cells. This supported the hypothesis of the original paper: “The effectiveness of antimalarial treatment against malaria will decrease with time as the malaria becomes resistant”.

The second hypothesis being tested by the malaria model was also supported by the results gathered. “Each subsequent application of a new strain of antimalarial treatment will have a lower overall effectiveness than prior applications due to increased malaria resistance”. There was an exhibited decline in the overall effectiveness of subsequent applications of antimalarial drugs in conjunction to the malaria cells becoming generally more resistant to the treatments, regardless of their strain.

The human model similarly supported the hypotheses put forth that: “The frequency of application of new strains of antimalarial treatments has little long term effect on the population of malaria within an ecosystem”. The trend represented by the results shows how unless the frequency of introduction of new antimalarial treatment strains was at a level beyond that which can be reasonably expected within the case study scenario, new streams of antimalarial treatments have little long term effect on the population of malaria within the model.

4. Evaluation and Conclusion

Though a general trend is presented in the graphs produced from the malaria model, the team noted a certain level of changeability in the results. Executing the same model with the same parameters has been found to produce differentiated sets of results, with variations present within the models’ running time and the graphs’ profile. A certain level of unpredictability is expected as the starting conditions for the model are randomised during the setup function.

An additional potential flaw with both models lies within the matcher, which is the part of the model that managed which treatment applications are successful by comparing the genome of the Malaia strain with that of the antimalarial drug. Due to the team being unable to replicate the breeding methodology put forth by Anabela Simões and Ernesto Costa (2000) within the allotted time, the portions of code that would have been responsible for representing an accurate breeding system within the model could not be used to handle accurate genome matching either. This meant the models had to fall back on a simple similarity slider. Without an accurate method of dividing which treatment applications are effective, the overall quality of the results is lower than what would have been desired.

The team has identified a number of changes they would have made to the projects had there been additional time available. One such example is the `breeding.nls` file developed by the team which includes multiple methods for genome mutation. The original plan was for the team to implement the full model put forth by Anabela Simões and Ernesto Costa (2000), however, due to time restrictions a suitable implementation was not able to be produced as mentioned in section 2.3 of this report. Instead, the team

was able to implement function “simple-swap-set-length-breed” which produces a similar effect to the methodology represented in the aforementioned paper with reduced effect.

The malaria model overall brought forth some interesting emergent properties and was a well rounded system that demonstrated the hypothesis put forward by the original source paper. However, one shortcoming of the malaria model and potential cause for some erratic results is the fact that the model represents an entire drug cycle taking place within the space of only a few ticks; experimentation indicates that the best results from the model can be seen when provided with a 20 tick treatment cycle. It is theorised that this is not enough time to allow the model to effectively represent an entire application cycle of an antimalarial drug.

The original study continued for 9 years before the first new treatment was introduced. This means that 1 tick in the model typically equates to 5.4 months in real time. It has been theorised by the team that additional emergent properties would have been displayed had the time scale of the malaria model been increased. These changes would have also provided a greater resolution of data with which to work with when producing the results section of this report and increasing the quality of graphs produced.

Again, looking at the malaria model, the transition from a Junior cell to a Senior cell is met when the junior cell has aged by 1 tick. A change that would improve this model would be to allow variation of the age in which the cell can transition into a Senior state, instead of hard coding the value at 1. This would allow for the possibility of delaying the transition into the Senior state thus delaying the reproduction process possibly leading to an emergent property in this area.

The team found the logging library was useful not only for diagnosing entity behaviour within the NetLogo models, but also helped with production of the first graphs as the model was being developed. However, because of the nature of BehaviourSpace, the logging library could not be used to generate the statistical data the team needed to generate their statistical data as each parallel instance of the model would try to write to the same log file. As BehaviourSpace supports its own data generation, it was later discovered that this was not an issue.

It can be concluded that the results of this paper support the hypotheses put forth by the source paper. The results also conclude that the introduction of subsequent treatments is less effective when applied to a malaria culture that has become resistant to a previous treatment. This paper also concludes that the frequency of introduction of new treatments is largely ineffective at controlling the long term spread of malaria within a population.

5. Acknowledgements

This template is heavily based on the PPIG template specified by Alan Blackwell. Many thanks to Alan for allowing us to use and modify this. Thanks also to Eleonora Bilotta, Thomas Green and Paola Kathuria for their help in defining and testing the original template.

The team would also like to acknowledge the guidance and assistance provided by Dr Ian Wood, Dr Simon Lynch and Ms Elisabeth Yaneske which made this paper possible.

6. References

Nsanzabana, C. et al. (2010) Quantifying the Evolution and Impact of Antimalarial Drug Resistance: Drug Use, Spread of Resistance, and Drug Failure over a 12-Year Period in Papua New Guinea. *The Journal of Infectious Diseases*, 201(3), 435–443.

Engelstädter, J. (2017) Asexual but Not Clonal: Evolutionary Processes in Automictic Populations. *Genetics*, 206(2), 993-1009.

Simões A, Costa E. (2000) Using Genetic Algorithms with Asexual Transposition. *Proceedings of the Genetic and Evolutionary Computation Conference, (GECCO '00)*.

World Health Organisation (2014) Malaria. World Health Organisation. [Online]. [17 November 2018]. Available from:

[<https://web.archive.org/web/20140903002027/http://www.who.int/mediacentre/factsheets/fs094/en/>]

Klein E.Y. (2013) Antimalarial drug resistance: a review of the biology and strategies to delay emergence and spread. *ScienceDirect*, 41(4), 311-317.

White N.J. (2004) Antimalarial drug resistance. *The Journal of Clinical Investigation*, 113(8), 1084–1092.

7. Peer Marks

Bellow are the teams peer marks split for the paper section of the module.

Name	Student Number	Peer Mark
Aaron Walker	Q5045715	4
Adam Precious	Q5068888	4
Ryan O'Donnell	Q5273477	3
Ladislav Baran	Q5127950	2