Modelling Epidemics with Cellular Automata

5591917

April 12, 2026

Contents

Foreword

Hello to anyone who comes across this essay. Before we get started, I wish to clarify something regarding the length of the content of this essay. This is primarily to ensure to the markers that I have not written too much, but also to alleviate any concerns that a future first or second-year student may have on first glance. By my best calculations, the length of the main written content of this essay (that being the written content in Sections 1–5) is somewhere between 13.5 and 14.5 pages in length; this is certainly on the higher end of what is acceptable, though importantly not so long as to disqualify any of said content from being read and considered.

Incredulity towards this statement is neither surprising nor unwarranted, however, as my main content spreads over more than 15 pages, as suggested by the contents page and the reality of my essay. This discrepancy can be justified, as it is a result of several additions which do not come under written content. Chief among these include a number of captioned images, a number of epigraphs, and a large amount of white space from using the page break command. While the latter 2 are solely for aesthetic purposes, I believe them to be a positive addition to my essay, hence why I have included them. With that explained, we shall begin.

1 Introduction

> Humans did not make history —
> we played host.
>
> ---
>
> *Pathogenesis*
> Jonathan Kennedy

In his book *Pathogenesis*, Jonathan Kennedy boldly makes the claim that infectious diseases, rather than human actions alone, have played the primary role in shaping human civilisation. This point undoubtedly has merit, and one consequence of this is that quickly and accurately modelling diseases and their spread is of the utmost importance. This endeavour is also highly technical, and many considerations have to be made when making appropriate models.

This first step in almost all models is to compartmentalise the host population into distinct groups with the various properties which we wish to model (Susceptible, Infected, Recovered, etc.). Successful models then have to describe how the population as a whole moves between these various compartments. One common way of doing this is by using systems of Ordinary Differential Equations (ODEs). In spite of the fact that the human population is inherently discrete, and the continuous nature of ODEs, David

Earn notes that for a large enough host population each compartment can be modelled as a continuous variable [1, p. 7], and thus ODEs can be used to link the compartments without issue.

The prototypical system used is the well-known SIR model [7, p. 612]

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -rSI,$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = rSI - aI,$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = aI,$$

where $r > 0$ is the infection rate and $a > 0$ the removal rate of infected individuals. While this model is undoubtedly useful, and shows up in literature dating back to at least 1927 [5, p. 713], there are some associated shortcomings. Most notably, this model is solely time-based; large epidemics tend to travel vast distances, making their spatial aspect increasingly important to consider as their scale increases. There are numerous ways to construct models that take into account both time and space; in this essay I will primarily argue for the utility of Cellular Automata, which from here onwards will often be referred to as CA, in both the singular and plural.

I will first outline the history and usefulness of CA, and describe their defining components. Following this, I will discuss creation of various 'useful' CA of progressively higher complexity and quality; initially the model will be deterministic and largely hypothetical. It will subsequently progress to being a stochastic SIRS model (Susceptible-Infected-Recovered-Susceptible) [3, p. 413], and be compared to alternative modelling methods. Following this, I will then discuss a variety of further improvements to the model, and how these improvements might be implemented and applied.

2 Creating Cellular Automata

2.1 What is a Cellular Automaton?

> This month we consider Conway's latest brainchild, a fantastic solitaire pastime he calls "life."
>
> *Scientific American, Vol. 223*
> [2, p. 120]
> Martin Gardner

Cellular Automata are a fairly recent mathematical invention, first tangentially discussed in the 1940s by Stanisław Ulam and John von Neumann. They then slowly gained notoriety throughout the 1960s, and first reached the wider scientific world in October 1970, when John Conway's Game of Life was published in *Scientific American*. They have since been used in a variety of scientific fields, including but not limited to epidemic modelling.

2.1.1 The components of a Cellular Automaton

Hadeler and Müller state [3, p. 19] that any CA has 4 defining components. Within [3, § 2], Hadeler and Müller describe each of these in far more detail than is needed for this essay; a brief definition of each concept will be given below, along with the location of the full definition and some extra notes regarding important subtleties.

Definition 2.1. [3, pp. 19–26] The *grid* is the discrete space over which the CA is run. A grid can be imagined as graphs with symmetries described in terms of a group [3, p. 19].

These grids are most frequently rectangular and abelian. The former allows for one-to-one projection onto standard pixelated screens, while the latter property helps standardise infection spreading; the chance of an infected individual spreading the disease in any valid direction should be equal, and if every individual is identical then the interactions between neighbours become commutative. However, neither of these conditions need to hold for a CA to function.

Definition 2.2. [3, pp. 26–28] The *neighbourhood* of a cell within the CA is a finite set of vertices which are used when determining the future state of a cell.

Perhaps surprisingly, the cells within the neighbourhood of a given cell $\mathbf{C}$ need not be connected to $\mathbf{C}$, or even include $\mathbf{C}$, though this is highly abstract and not necessary to discuss in further detail within this essay. For

regular grids, such as the ones which we will use, certain neighbourhoods can be defined by using various norms [3, p. 26] (more on these later).

Definition 2.3. [3, pp. 28–30] The *states of a single cell* are admittedly self-descriptive; there exists a finite set of elementary states **E**, and functions are defined on the graph mapping each cell of the graph to one of these states [3, p. 28].

Definition 2.4. [3, pp. 30–32] The *local function* is the function that determines the state of a cell **C** at a timestep *n+1* based on the states of the cells in the neighbourhood of **C** at the timestep *n*.

This is the key mechanism governing the evolution of a CA.

2.1.2 Why are we using Cellular Automata?

At first glance it may seem rather naive to use CA to model phenomena as complex as epidemics, given their simplicity and highly localised interactions. Indeed, these limitations are severe enough to the point where continuous methods of space-time modelling, such as Partial Differential Equations (PDEs), are often used; some examples relating to diseases such as the West Nile Virus can be found in [1, §. 8].

Hadeler and Müller mention this in [3, p. 1], noting the 'beautiful mathematical theories and powerful applications' associated with PDE models. However, their elegance comes at the cost of restrictions imposed by the regularity and existence problems, which are typically unrelated to the physical model which they attempt to describe. This does not stop PDEs from being incredibly useful; indeed, numerical methods such as finite-difference and finite-volume can be used on various models to output results from which many powerful conclusions can be derived.

In contrast, CA models are discrete and governed by very simple rules, at least initially; these factors are actually their greatest strengths. Due to their nature it is almost always feasible to run computer-aided simulations and analyse the data which results. In the case of small to medium-sized 2-dimensional CA, their simulations are able to be viewed in their entirety on a computer screen.

Although this is computationally expensive, especially in comparison to similar ODE systems, the ability to see the disease spread through a physically arranged population is a unique benefit of using CA, which can be further enhanced by modifying the original state of the system. By limiting initial infections to a small corner of the system, one is able to clearly observe travelling fronts of infection, which offers powerful insights into how to effectively implement measures such as quarantine zones.

Now that we have defined the key aspects of a CA, and justified a potential niche for them, it makes sense to construct one.

2.2 Creating a deterministic Cellular Automaton

By choosing specific versions of the above 4 concepts, any number of CA can be constructed. The most famous of these is undoubtedly Conway's *Game of Life*, whose rules can be found in the Gardner article [2, p. 120], and happens to be entirely deterministic. Importantly for us, by basing our selections on epidemiological concepts, we should be able to construct a CA which communicates some realistic concepts of epidemics.

The 1996 Johansen paper will provide much of the basis for our first model, as the rules that the model epidemic will follow are very clearly laid out [4, p. 46], as well what the modifications are for the deterministic epidemic [4, p. 48]. The rules that govern the evolution of the system over a single timestep are listed verbatim as follows:

1. The disease spreads to nearest neighbours (*n.n.*) of susceptible individuals with probability q.
2. Sites with infectious individuals become inactive (or recovered).
3. New susceptible individuals are introduced to inactive sites according to a specified probability distribution for the inactive period.

   Remark. Notably this allows for re-infections of individual cells.
4. Susceptible individuals become infected with a probability c. We will refer to this as indirect transmission.

These points, in combination with the modifications and observations gleaned from [4, pp. 46–48], amount to the following definition of our deterministic CA:

Defining Our Cellular Automaton

The grids used in our automata will be square lattices with side length of 100 cells, resulting in the model representing a population of 10,000. This is well within the range of grid sizes described in [4, p. 47].

The neighbourhood used is the *von Neumann neighbourhood* of range 1, the general case of which for the 2-D grid is defined below:

Definition 2.5. The *von Neumann neighbourhood* of range r for a given cell (x_0, y_0) in the 2-D grid is defined as the cells with Manhattan norm $\leq r$, i.e.

$$N^v_{(x_0,y_0)} = \{(x, y) : |x - x_0| + |y - y_0| \leq r\}$$

Johansen mentions that the Moore neighbourhood, defined directly below, can also be used without issue [4, p. 47].

Definition 2.6. The *Moore neighbourhood* of range r for a given cell (x_0, y_0) in the 2-D grid is defined as the cells with Infinity norm $\leq r$, i.e.

$$N^m_{(x_0,y_0)} = \{(x, y) : \sup\{|x - x_0|, |y - y_0|\} \leq r\}$$

I have chosen to exclusively use the von Neumann neighbourhood, not only for ease of use, but also due to the fact that Figure 4.(a) [4, p. 48] appears to have been generated using the von Neumann neighbourhood.

The states that a cell can take are Susceptible, Infected, and Recovered, much like for the ODE model mentioned in [7, p. 612]. By setting their colours differently in our code, we will be able to observe the behaviour of each cell in detail. This is a notable change from [4], where only black and white are used to draw the diagrams. We deviate from this not only for aesthetic reasons, but also to easily distinguish between susceptible and recovered individuals.

As there are 3 types of cells, we shall define a local rule for each of them, though only one of them technically uses the neighbourhood described above.

1. For a susceptible cell C at timestep N, C will become infected at timestep $N+1$ if an infected cell is in the von Neumann neighbourhood of range 1 of C at timestep N. If no infected cell is within the neighbourhood of C at timestep N, C will remain susceptible.

   Remark. This means that $q = 1$ and $c = 0$, in the context of [4, p. 46]

2. For an infected cell C at timestep N, C will become removed at timestep $N+1$.

3. For a cell C which was infected at timestep N-1, and by extension 'freshly' became removed at timestep N, C will become susceptible at timestep $N + t$. This t is kept deliberately vague, as it can easily be adapted to be any constant value, or alternatively a statistical distribution. For this CA, as described in [4, p. 48], $t = 35$.

   Remark. Similarly, we can say that for an infected cell, $t = 1$.

2.3 Observations and Limitations

The above deterministic CA is capable of providing some interesting insights. Most notably, by limiting the number of infected individuals at $t = 0$, travelling fronts of infected individuals can clearly be observed, as well as the annihilation which occurs when two fronts travelling toward each other collide. Images of these phenomena can be found within Johansen's paper [4,

p. 48]. However, there exist glaring shortcomings with the perfectly deterministic model, which must be fixed for the model to convey any sense of realism.

These issues have come as a direct consequence of making the model perfectly deterministic, and are not at all surprising; there will naturally be a lack of realism that comes with the ease of concept and programming. More specifically, the variables $q = 1$, $c = 0$, $t_{I \to R} = 1$, and $t_{R \to S} = 35$ combine to make behaviour highly unrealistic. By addressing these issues, we will improve our ability to use our model to usefully make observations and recommendations in the context of actual epidemics.

3 Increasing realism

> Continuous improvement is
> better than delayed perfection.
>
> Mark Twain

3.1 Making a Cellular Automaton stochastic

As highlighted above, the deterministic nature of numerous defining variables result in our current model having very limited utility. Thankfully, all 4 of these variables can be easily modified in a way to make our model more stochastic, and by extension more realistic:

1. Making $0 < q < 1$ would symbolise a base level of immunity in the human population, either via natural means or through vaccination.

2. Making $0 < c \ll 1$ would symbolise a degree of indirect transmission, for example contracting the flu through exposure to lingering droplets in public spaces.

   Remark. In the context of epidemics, $c \ll q$, as indirect transmission should always be far less effective at transmitting the disease than direct transmission.

3. Modifying the times taken to for individuals to move from $I \to R$ and $R \to S$ from a constant value to a function that can output a range of values.

Johansen uses multiple of these ideas at various points throughout [4]. On [4, p.47], the time taken for cells to move from $R \to S$ is modified to be 'exponentially distributed', on [4, p.49] the above time period is left deterministic but a value for $c = 0.015$ is used, and then Johansen goes on to vary $q < 1$ on [4, p.50], to show the effects of immunisation. These have a wide variety of useful insights, though replicating them is of little interest to us, given the fact that parts of these models remains deterministic. In the next part of our essay, we will create a more accurate model by altering all of the above variables, in an attempt to accurately model COVID-19.

3.2 Attempting to model a real disease (COVID-19)

Up to this point, our CA has only been theoretically discussed in the context of modelling caricatures of epidemics, where specific variables have been kept as deterministic in order to highlight specific properties, such as travelling fronts and their organisation, even in systems where indirect infection can occur to any susceptible cell at any point. The next step, and the crux of this essay, is to use a Cellular Automaton to try to model a realistic epidemic.

In order to do this, we will need to alter Johansen's models even further and make every single parameter reflective of the stochastic nature of real epidemics. Due to the exceedingly large data samples available, COVID-19 is the disease of choice to be modelled, with articles such as [6] and [8] containing values for some of the parameters. The values below will describe the rules of the epidemic which we wish to follow:

1. The probability that a susceptible individual will become infected upon contact with an infected individual at each timestep, $q = 0.209$. Taken from [6, p. 63].

2. The probability that an infected individual recovers at each timestep, $p_{I\to R} = 0.0575$ (If this transferral from $I \to R$ were instead implemented deterministically for this data, we would have that $t_{R\to S} \approx 17$). Also taken from [6, p. 63].

   Upon first glance, it may seem unusual that we have taken values from a model system of ODEs as the variables of our CA; this is easily justified, however, as [1, p. 7] states that a generalisation is made when moving from a discrete system into a system of ODEs. If the parameters are valid for the ODE system, they must inherently be valid for the initial discrete system, meaning that we can use them withoout issue.

3. The probability that a recovered individual will become susceptible to re-infection, $P_{R\to S} = (T_{R\to S})^{-1}$. From [8, §. 6], Shen, Chen, Li, Huang, and Ma found a wide variety of mean re-infection times, with most being a little under a year. Taking $T_{R\to S} = 350$ gives us the value $P_{R\to S} \approx 0.00286$, which is the value we shall use.

4. The probability that a susceptible individual becomes infected due to indirect causes at each timestep, c. Unlike the other interactions, c does not make sense to be constant; as such, its correct modelling is worthy of further discussion, as seen below.

Modelling c

Due to the lack of literature on the topic, the burden falls on us to find an appropriate value or function that describes c. We assume that c scales linearly with the proportion of the population that is infected at any point, thus giving us an interim model of $c = \alpha \times I \div N$, where α is a constant to be determined.

According to the Office of National Statistics[4], the week ending on March

[4]https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/bulletins/coronaviruscovid19infectionsurveypilot/25march2022, last accessed on 24/02/2026.

19th 2022 had the largest proportion of the UK population infected with COVID-19, with approximately 1/16 of the population testing positive. Our function for c would then take a sensible value of our choosing when $I \div N = 1/16$.

For [4, p. 50], Johansen takes $c = 0.015$; I personally believe that this is too high to be a realistic value at any point in the disease progression. As we are assuming that $I \div N \leq 1/16$ at all times, we then make the arbitrary decision that $c = 1$ when $I \div N = 1/16$. Solving for α gives $\alpha = 0.16$. Therefore our function for c reads:

$$c = 0.16 \times \frac{I}{N}.$$

3.3 Our results

Our next step is to run simulations with various versions of the optimised model, that is we will run one of the model without indirect infection, and then one with it. While the former may seem pointless, there is a need for it; the code used in the models means that the model **with** indirect infection includes all of the code used in the model **without** it.

Therefore, simulating the 'interim' model will allow us to confirm that the programming of the direct infection works as intended. It will also allow us to see it operate in a vacuum, meaning that the travelling fronts will be easier to analyse; there will only be 1 travelling front for the first few hundred timesteps.

Both models shall be simulated over a 100 by 100 grid, giving a population of 10,000 individuals. Both models will also have a partially determined initial state; only the 100 individuals in the upper-leftmost 10 by 10 square will be able to be infected, with infection occurring at a rate of one quarter. This choice is made to be able to see a large travelling front, especially in the interim model's case.

By simulating, we will also be able to analyse how useful our models are, both in theory and in reality. The first of these analyses will be done by comparing these stochastic simulations to corresponding systems of ODEs, which will be stated within the sections. Additionally, the model with indirect infection will be compared to relevant real-world data.

3.3.1 Interim model: Stochastic model without indirect infection

As suggested by the title of this section, the interim model does not convey indirect infection; however, all other key parts of the epidemic modelling are included. An image of this model at an interesting point is displayed below as Figure 1.

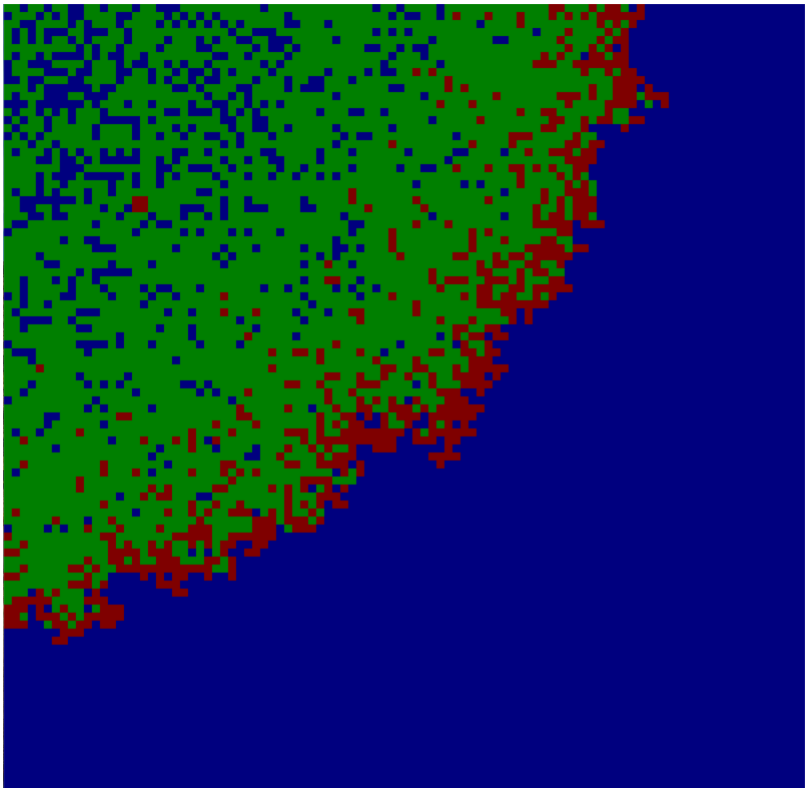

Figure 1: Snapshot of the spatial evolution of the interim model, taken after 200 timesteps had elapsed.

As we can see from Figure 1, a travelling front spanning across the entire effective width of the population is clearly visible. In comparison to the deterministic version, this front is both much thicker and notably less homogeneous; this trivially makes sense given the stochastic nature of this model, along with the relatively low probability of a given infected cell recovering quickly. We also observe patches of infected individuals surviving behind the travelling front, which is a natural result of the above.

The ODE system which directly corresponds to this model is shown below

$$\begin{aligned}
\frac{\mathrm{d}S}{\mathrm{d}t} &= 0.00286 \times R - 0.209 \times S \times I, \\
\frac{\mathrm{d}I}{\mathrm{d}t} &= 0.209 \times S \times I - 0.0575 \times I, \\
\frac{\mathrm{d}R}{\mathrm{d}t} &= 0.0575 \times I - 0.00286 \times R,
\end{aligned}$$

and comparing this model with our stochastic one has some very interesting results, which are displayed in Figure 2.

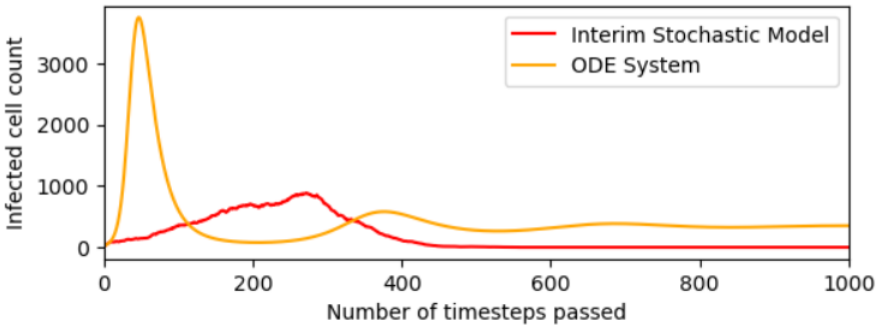

Figure 2: Time-evolution of the number of infected individuals within the interim stochastic model, and the above ODE system, which mirrors its parameters.

Firstly, the peak number of infections is far lower in our stochastic model, which only reaches a peak of 800 simultaneous infections in comparison to the ODE system's peak of 3700. Second, the ODE system eventually reaches an equilibrium state with about 1 in every 30 individuals being infected, whereas the stochastic model dies out completely after roughly 500 timesteps. Both of these phenomena can be explained by the fact that we have introduced a spatial element to the stochastic model.

With infections being exclusively dependent on direct infected neighbours, the stochastic model cannot infect nearly as many individuals in any one timestep as the ODE system can, which peaks within 50 timesteps. For reference, the stochastic model tends to peak after more than 250 timesteps.

Regarding the second point, for the infection to not die out, infected individuals must stay infected for sufficiently long to interact with susceptible individuals 'behind' the travelling front. This occurs either due to some infected individuals taking an exceedingly long time to recover, or due to recovered individuals behind the travelling front quickly becoming susceptible. With the variables which we currently have, the chances of this happening in the stochastic model are extremely low, meaning that the vast majority of simulations experience extinction of infected individuals after 500 timesteps. Thankfully, these flaws are largely irrelevant; we have shown that the stochastic model might work in theory, and as such we now observe our 'true' model.

3.3.2 Our stochastic model of COVID-19

Our 'true' model is exactly the same as above, but with the addition of indirect infection as a phenomenon; this, rather ironically, makes pictographic observations rather less entertaining, as the behaviour rapidly homogenises

after approximately 100 timesteps. As such, no images of the model's evolution are included in this essay, though this shall not detract from the analysis of the model.

The addition of indirect infection also changes the associated ODE system, and must be accounted for. Thankfully, this adjustment is simple to implement: indirect infection affects every single susceptible cell in the global system, thus the number of cells expected to be infected at each timestep is $(0.16 \times I/N) \times S$. As this is functionally identical to the other part of the ODE (N is just a constant), it can be absorbed; the system then reads as

$$\frac{\mathrm{d}S}{\mathrm{d}t} = 0.00286 \times R - 0.369 \times S \times I,$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = 0.369 \times S \times I - 0.0575 \times I,$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = 0.0575 \times I - 0.00286 \times R.$$

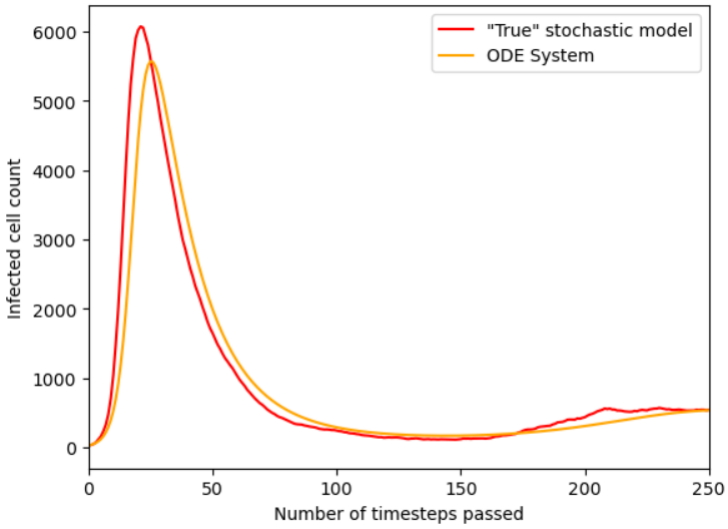

Figure 3: Evolution over the first 250 timesteps of the number of infected individuals within the improved stochastic model, and the improved ODE system.

Comparing our 2 models yields some very reassuring results, as shown in Figures 3 and 4. Most importantly, the behaviour of both models matches very closely, both in the short and also, more importantly, the long-term. The system of ODEs eventually reaches a steady state where infected individuals make up (almost exactly) 4% of the population; while the stochastic

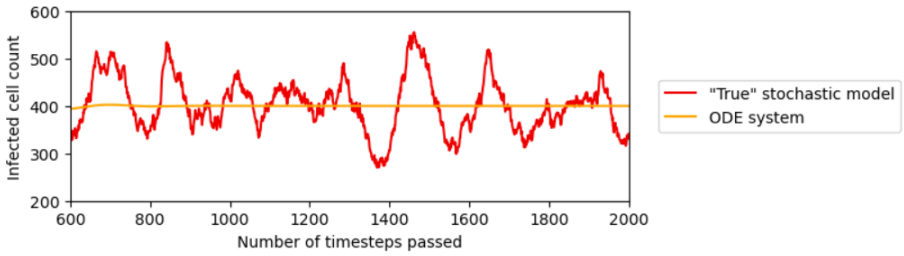

Figure 4: Evolution from 600–2000 timesteps of the number of infected individuals within the improved stochastic model, and the improved ODE system.

model is never 'steady', the random walk shown by the infected population appears to centre around a value very close to 4% of the population. At the bare minimum, this shows that CA may be viably used to model infectious diseases, such as COVID-19.

However, there is a **very** big flaw with our CA, in that it completely breaks one of the key assumptions we made in Section 3.2. When modelling c, we worked on the basis that the infected population of the model would peak at 1/16 of the total population, and created our formula based on that fact. In reality, the infected population within this model peaks at over **60%**, nearly 10 times higher than our ideal peak.

While it might be true that the c-value chosen when modelling c was too high, the discrepancies found cannot be solved by simply modifying the indirect infection formula; optimisation may help a little, but even our interim model had a peak in excess of 800, higher than our real-world data suggests. Therefore, further improvements **must** be made to the model for it to convey a greater sense of realism.

4 Further improvements

> Remember: things can be bad,
> and getting better.
>
> *Factfulness*
> Hans Rosling

As mentioned above, our model of COVID-19 has many areas in which it needs to be improved. There are also many methods of improving the model; some of them attempt to directly tackle we issue that we have addressed at the end of the previous section, while others 'merely' smooth out a few edges by introducing more generic concepts. A few of these methods are listed below, along with reasoning behind their viability and hypothetical methods for their implementation:

1. Immunity

   While immunity is already modelled to some extent by the existence of the recovered category, its current implementation fails to convey the concept of partially faded immunity, and the transfer mechanism from recovered to susceptible leads to some individuals losing immunity far more rapidly than they should

   These two problems have two rather different solutions. The former problem could hypothetically be solved by introducing a number of partially susceptible states, i.e. each of the k steps increases the probability of infection from nearest neighbours by $\frac{0.209}{k}$, until the fully susceptible state, at which point the chance of infection is as in our original model.

   In contrast, the second problem has an easier fix; we can switch the mechanism used to a normal distribution, with mean 350 and a suitable standard deviation which we would suggest, as [8] does not provide us with useful insight. These fixes can even be combined into one, which aims to achieve both simultaneously; this may be done by using a number of normal distributions with means $\frac{350}{k}$ and a sensible standard deviation.

2. Quarantines & Lockdowns

   Quarantines and Lockdowns, as shown during the real COVID-19 pandemic and various follow-up waves, were extreme but effective measures in controlling the disease's spread. As such, being able to implement various forms of them into our model would give us a good idea of how strictly they need to be implemented in 'real life' situations.

Quarantines would act on individual cells, and could work by dropping an infected cell's chance of infecting its neighbours significantly (or to zero) for an arbitrary period of time. Alternatively, they could work by dropping the chance to zero until recovery, but not have a perfect chance of being implemented optimally, i.e. an infected cell evades the screening process. Lockdowns would affect much larger areas of the grid, if not the entire grid, and would function similarly to quarantines on every infected cell within them.

3. A collection of smaller CA

   This idea is undoubtedly more abstract than the ones above; nevertheless, it ought to be mentioned as a potential modification of a coupled patch model [1, §. 7]. Within this hypothetical model, a large number of smaller CA will represent pockets of a given population; notably they need not all be the same size. They are then all able to act independently, and inter-CA interaction will be limited to migration events, where 2 cells in 2 different CA swap states.

   Interestingly, this can be done in a variety of ways; either every CA has a chance to interact with each other, or the CA are arranged within a more concrete network, and only the CA which are 'connected' can interact with one another. This could be determined by an arbitrary graph, or even another CA, where each cell is home to one of our smaller CA. This is undoubtedly the hardest version to model, but it is very much worth postulating, given the primary focus of this essay, and the application it could have towards modelling disease migration between towns/cities.

5 Conclusion

> Everything has to come to an end, sometime.
>
> Tip, *The Marvelous Land of Oz*
> L. Frank Baum

Reiterating what was said in the introduction, modelling epidemics is an extremely important task, and many models have been created and applied over the past 100 years. These models have various advantages and disadvantages, and this essay discusses those associated with Cellular Automata. Over the course of this essay, we have motivated the reasons for using CA, their fundamental components, and what aspects of a disease would need to be communicated by them in order for our CA to be useful models.

Following this, we have created a model CA, and compared it to an ODE system with similar parameters and starting conditions. We then discussed further alterations which can feasibly be made to our current model to convey a greater sense of realism. This creative process has clearly demonstrated some of the most consequential advantages and disadvantages of using CA.

During my work on this essay, the largest problem I encountered with using CA in this context was the cost associated with computation. Not only did my CA models take significantly longer to run than the corresponding ODE models, their creation involved a very large struggle on my end. Needless to say, I was very much out of my depth during most of the process; even the models created in Section 3 took a monumental effort, as simple as they may be in reality. This is also the reason why the alterations in Section 4 are merely suggested instead of formally implemented, as my coding knowledge is not sufficient to implement them properly.

The fact that they can be conceptualised, however, is a by-product of the greatest positive trait of CA that I have demonstrated: their ability to communicate both temporal and spatial traits effectively. The fact that the long-term behaviour of my 'best' CA matches with its corresponding system of ODEs is highly reassuring, as it means that temporal traits match well.

Additionally, the further adaptations suggested in Section 4 are more suited to CA than ODE systems; some, such as sequential immunity loss, are possible but complex; a 10-step immunity loss scheme would lead to a system of 12 ODEs, for example. Even more extremely, the implementation of lockdowns would be entirely impossible on the ODE system, as there is no concept of 'space', and by extension nowhere to lock down.

If I had, rather ironically, more time or space to continue this essay, there is a quite solid path which I would likely end up following. Firstly, I would (attempt to) implement some or all of the ideas mentioned in Section 4, either in isolation—such as how Johansen did with his adaptations in [4, pp. 47–50]—or in tandem, such as how we did in Section 3.2. My methods of comparison would also change; I would outline and construct a relevant system of PDEs, which I would then use as a comparison. As this is largely hypothetical, I have a very limited idea of how to do this, though from a cursory glance the content in [1, §. 8] would likely provide a good start.

And with that, I have nothing else to say; this essay is now finished. I hope you enjoyed reading it as much as I enjoyed writing it.

References

[1] Fred Brauer, Pauline Van den Driessche, Jianhong Wu, and Linda JS Allen. *Mathematical epidemiology*, volume 1945. Springer, 2008. URL https://link.springer.com/book/10.1007/978-3-540-78911-6.

[2] Martin Gardner. Mathematical games. *Scientific American*, 223(4):120–123, 1970. ISSN 00368733, 19467087. URL http://www.jstor.org/stable/24927642.

[3] Karl-Peter Hadeler and Johannes Müller. *Cellular automata: analysis and applications*. Springer, 2017. URL https://link.springer.com/book/10.1007/978-3-319-53043-7.

[4] Anders Johansen. A simple model of recurrent epidemics. *Journal of Theoretical Biology*, 178:45–51, 1996. doi: https://doi.org/10.1006/jtbi.1996.0005. URL https://www.sciencedirect.com/science/article/pii/S0022519396900053.

[5] William Ogilvy Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, 1927. URL https://royalsocietypublishing.org/rspa/article-pdf/115/772/700/24858/rspa.1927.0118.pdf.

[6] Loukas Kyriakidis and Michail Kyriakidis. Self-consistent estimation of ordinary differential equation parameters describing dynamical systems: A case study of COVID-19 in Germany. *WSEAS Transactions on Biology and Biomedicine*, 22:53–66, 2024. URL https://elib.dlr.de/207780/1/Self-consistent%20estimation%20of%20ODEs%27%20parameters.pdf.

[7] J. D. Murray. *Mathematical Biology*. Springer-Verlag, New York, 2nd corrected edition, 1993. URL https://link.springer.com/book/10.1007/978-3-662-08542-4.

[8] Hongwei Shen, Dingqiang Chen, Chenglin Li, Tingting Huang, and Wen Ma. A mini review of reinfection with the sars-cov-2 omicron variant. *Health Science Reports*, 7(4):e2016, 2024. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC11007061/#hsr22016-sec-0120.

A My little Odyssey: Coding the Model

> Be strong, saith my heart; I am a soldier; I have seen worse sights than this.
>
> Odysseus, *The Odyssey*
> Homer

As mentioned regularly in sections 3, 4, and 5, my coding skills coming into this project were rather unsatisfactory. This initially gave me great problems, as I had already invested a significant amount of time and effort into finding appropriate literature for the project, and my ideal essay revolved around being able to create a model. This concern was further exacerbated by my lack of confidence regarding coding; even if I could conceptualise what I wanted my code to do, putting it into programming was a daunting task, in spite of how simple the code might appear to an outsider.

Thankfully, I was able to locate a very useful page on GitHub, which contained a number of interesting models, including the one I ended up adapting for my own work. The page in question can by accessed through https://github.com/temp3rr0r/CellularAutomataEpidemicModels; the specific file I used is found by going into the 'Working Models' folder and then selecting GameOfLife2DSIRSImage.py.

This model already had a number of highly useful features that I wanted to keep. For example, the model had variables for the number of cells on the x and y-axes already built-in, meaning that I could change the size of the population quickly and prove concepts and test variables on smaller grids, before using them on larger ones. However, some of them proved largely vestigial in the context of my model; the hexagonal modification is a prime example of this, which I swiftly removed, mostly for aesthetic reasons.

In spite of its many positives, the model was missing some mechanisms that I wanted to use. Critically, the state of the model at $t = 0$ was governed by a function which randomly distributed infected cells over the entire grid, at a rate determined by the programmer. While useful, I wanted to showcase the phenomenon of travelling fronts; this ideally required only a small corner of the grid to be infected at $t = 0$, which would allow for clear propagation.

With much help from another student in my tutor group, this change was successfully implemented, allowing me to initiate the infection at any specific location within the grid I wish. From there, I attempted to add the phenomenon of indirect infection, which involved creating a further variable, δ, which changed depending on the number of infected individuals, as well as the size of the initial population.

Thankfully, this turned out to be rather easy to do; as the original code already output graphs, it was necessary for it to keep a count of the infected cells at every timestep; the creation of our desired variable was then trivial. Its implementation into the model was also relatively simple; by using a basic 'elif' statement in the infection rules allowed us to introduce it after direct infection had taken place.