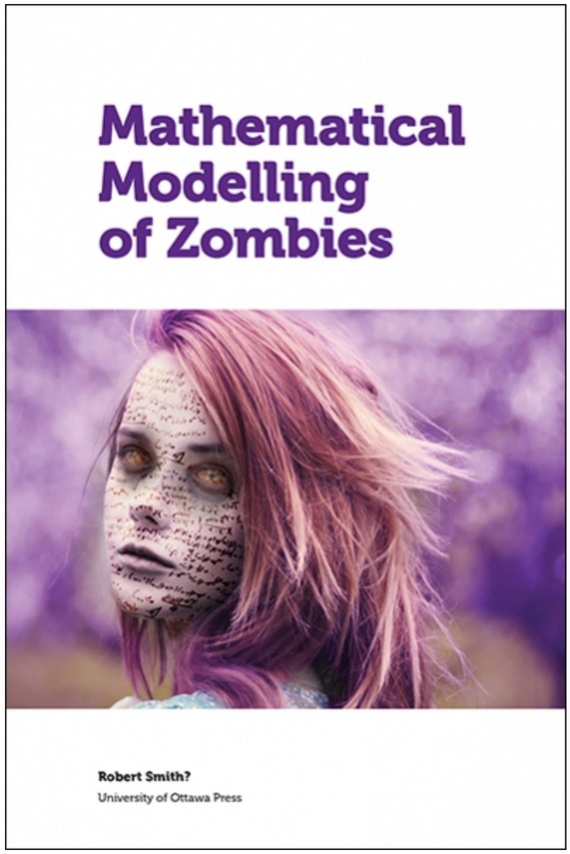

THE SOCIAL ZOMBIE: MODELLING UNDEAD OUTBREAKS ON SOCIAL NETWORKS

Laurent Hébert-Dufresne, Vincent Marceau, Pierre-André Noël, Antoine Allard, and Louis J. Dubé

Département de physique, de génie physique, et d'optique
Université Laval, Québec (Québec), Canada G1V 0A6

Hébert-Dufresne L., Marceau V., Noël P.-A., Allard, A., and Dubé L.J., *The Social Zombie: Modelling the undead outbreaks on social networks*, chap 9., pp 149-170, in Mathematical Modelling of Zombies, ed. Robert Smith? (University of Ottawa Press, 2014)

THE UNDEAD

authors of THE SOCIAL ZOMBIE

Laurent Hébert-Dufresne, Vincent Marceau, Pierre-André Noël,
Antoine Allard, and Louis J. Dubé

Louis J. Dubé received his B.Sc. in physics from Université Laval (1973) and completed his graduate studies at Yale University (M.Sc. 1974, Ph.D. 1977). This was followed by years of directed random walks, first to Oxford, UK (1978-79) and then to Freiburg, Germany (1979-86) where he emerged with the German doctorate and the Habilitation (Dr.rer.nat.habil. 1986). He joined at this point the professorial ranks at his *alma mater* Université Laval; this was the end of the first cycle. He has so far resisted other basins of attraction (Bordeaux, 1992-93; Paris, 2002- 2008) only to return to the original attractor. He had the chance to reinvent himself several times from collision theory to complex dynamical systems. He embodies a *RSIR* model: he was born *ready* for, quite *susceptible* to, rapidly *infected* by and still not *recovered* from scientific research and education. In the course of this work, he learned that zombies were, by strict definition, not vampires.

Pierre-André Noël , Antoine Allard, Laurent Hébert-Dufresne and **Vincent Marceau** were born and subsequently raised in the Province of Québec. Unbeknownst to one another, they shared first hand zombie encounters early in life (PAN: The Evil Dead, AA: Saw a bunch of them kicked out of a bar once, LHD: Zombies Ate My Neighbors, WM: Michael Jackson's Thriller). Unprepared at the time, they ignored their calling until their paths intertwined at Université Laval where they received their B.Sc. in Physics (PAN: 2005, AA: 2006, LHD & WM: 2009). They eventually chose to face the trials and tribulations of graduate studies (purification by fire) within the same holy halls, where they gathered under the Dynamica flag of Prof. Louis J. Dubé's research group on the dynamics of networks (actually PAN started with mosquitos, but as they say: blood sucking to flesh eating is a one way trip). Undeterred and hardened by the growing challenges, they have prevailed against papers, degrees and diplomas (PAN: M.Sc. 2007, Ph.D. 2012; AA: M.Sc. 2008, Ph.D. 2014; LHD: M.Sc. 2011, Ph.D. 2014; VM: M.Sc. 2011). They have now returned to face the zombies of their childhood nightmares for the single greatest modelling challenge of all times: the undead hordes... to save us all, or go mad trying.

Scully: *Zombies are just projections of our own repressed cannibalistic and sexual fears and desires. They are who we fear that we are at heart. Just mindless automatons who can only kill and eat.*
Mulder: *Well, I got a new theory. I say that when zombies try to eat people, that's just the first stage. You see, they've just come back from being dead, so they're gonna do all the things they missed from when they were alive. So first, they're gonna eat. Then, they're gonna drink. Then, they're gonna dance and make love.*

— The X-Files, season 7, episode 19.

Abstract

According to Mulder's theory, the zombies will eventually fall on each other and make love. However, be it for love or evil, the cold hard reality remains that the actions of the undead, just as those of the living, are also structured by simple constraints of social or spatiotemporal nature. In this chapter, we improve upon the standard zombie outbreak model by considering the underlying social network of the living and the horde behaviour of the undead. This model is then further improved by considering the adaptive nature of social interactions: people usually tend to avoid contact with zombies. Doing so captures the coevolution of the human social network and of the zombie outbreak, which encourages humans to naturally barricade themselves in groups of survivors to better fight the undead menace. And then? Better stack goods, arm yourself and be patient, for the undead hordes are there to stay; hopefully dancing and making love.

Introduction

Whether it is going to work/school (roads and public transport networks), updating a Facebook status (internet, online social networks, world-wide web), meeting friends or getting a high-profile job (acquaintance networks), calling abroad or getting directions using a GPS (satellite networks), or simply turning on the radio (electrical and information networks), we keep encountering networks in our everyday lives. By improving our understanding of the structure of such networks, we are increasingly able to derive as much benefit as possible from the advantages that networks have to offer, while efficiently protecting ourselves from the disadvantages.

Being social animals, people interact with each other for a variety of reasons: friendships, family bonds, sexual partnerships, business relations and so on. At the population level, these interactions sum up to form a giant web: the *social network*. How individuals are connected to one another in the social network depends on the nature of the considered interactions. Within this structure, people may exchange or transmit information, opinions or infectious diseases, to name a few, while the underlying social network shapes the propagation dynamics.

Studying social networks has been quite useful in the recent past to understand and, to a certain extent, predict the propagation of infectious diseases in human populations. While research has mostly been focused on containing and fighting "traditional" emerging infectious diseases such as HIV and influenza, another threat has been grossly underestimated and is now imminent: the zombie apocalypse.

As the hands of the clock push us inexorably closer to the End Time, experts believe that preparation is now imperative to ensure any future for humanity. Using concepts from network theory and contact network epidemiology, we present in this chapter a novel approach that models zombie outbreaks and human counterattack actions. This model will help authorities to conceive

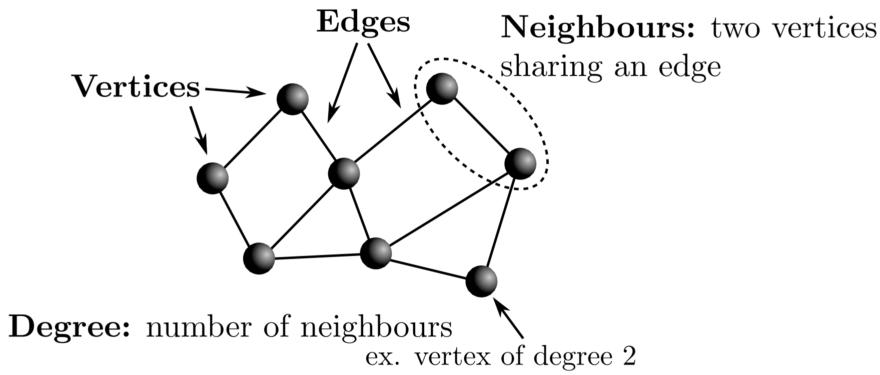

Figure 1: A graph of 8 vertices and 11 edges.

and test different resistance strategies beforehand, and be as prepared as one possibly can for *when there will be no more room in Hell and the dead will walk the Earth*[1].

1 Modelling a zombie invasion: a practical guide

To model a zombie invasion, we first need to introduce some basic notions of network theory. We will then define rules governing how people and zombies should interact.

1.1 Contact networks 101

A contact network is represented by a *graph*: a set of dots linked by lines. Dots and lines are generally referred to as *vertices* and *edges*, respectively, and we will use this nomenclature throughout this chapter. Figure 1 shows a simple example of a graph.

When modelling social networks, individuals are represented by vertices and two individuals are linked by an edge if they interact with each other. In the case of a zombie invasion, any kind of direct interaction may lead to infection. Therefore social interactions such as friendship, family bonds or school/workplace acquaintances are taken into account in the model.

Two vertices sharing an edge are said to be *neighbours* and the number of neighbours of a given vertex is called its *degree*. A property of real social networks that plays a major role in the disease propagation is their *degree distribution*, noted as $\{p_k\} = [p_0, p_1, p_2, \ldots]$. Each p_k gives the fraction of the vertices that have k neighbours, which is also the probability for a randomly selected vertex to be of degree k. For example, the degree distribution of the graph shown in Figure 1 is $\{p_k\} = [0, 0, 1/2, 1/4, 1/4, 0, 0, \ldots]$.

1.2 Building a contact network

The structure of a network contains *a lot* of information. In order to acquire this information through a survey, we would have to ask each vertex who their neighbours are. While this might be simple for a very small network, it is usually not possible for larger populations. Indeed, it would be much simpler to survey a small fraction of the population and then figure out what information this gives us about the whole network.

[1] Dawn of the Dead, George Romero, 1978

Now suppose that we have the results from a survey that asked 100 people how many neighbours they have (i.e., what is their degree). If the sole other information you have about the network is that there are a total of 10,000 people in it, how do you fill in the blanks? How do you obtain the degrees of the other 9,900 other vertices and how do you join all the vertices together?

The first question is easy to answer by requiring consistency: the degree of all the 10,000 vertices should be chosen such that if we would ask any randomly chosen 100 of these vertices to answer the same survey, the result should be similar to the one previously obtained. In practice, this is usually done by obtaining the degree distribution $\{p_k\}$ of those 100 people who answered the survey and then assume that the remaining 9,900 follow the same probability distribution. We can thus pick them one by one, assigning degree k to each new vertex with probability p_k.

The second question, how to join these vertices together, may be less intuitive. We could design a complicated process based on friendship, chances of encounters in the street, who is going to which supermarket... But if we do not have this information, should we make it up?

In all cases, *the best choice* is to use the little we do know, and nothing that we do not. In our case, this comes out as forcing vertices to have the degrees that we earlier chose while randomly assigning edges between them. One possible algorithm for doing this goes as follows: for each of the N vertices in the network, place in a bag k pieces of paper each bearing a tag uniquely identifying that vertex of degree k. Shuffle, draw two pieces of paper, assign an edge between the two corresponding vertices, destroy the drawn pieces of paper and repeat until the bag is empty. When it will be required to build a network from this information, we will use a procedure very close to the one just described. Of course, everything will be automatized in a computer program. The networks considered in the remaining of this chapter will be constructed by this algorithm and therefore simply defined by their size N and their degree distribution $\{p_k\}$.

1.3 The rules of the game

Until now, all the vertices in the network were intrinsically the same. This is clearly undesirable for our model since we do not expect a zombie to behave the same way as an healthy human. However, it is probably acceptable to say that all zombies behave as “a typical zombie”, and that all humans behave as “a typical human”.

We will differentiate individuals into three different states: those who have not been infected yet (denoted S for *survivors*), those who have been infected (denoted Z for *zombies*) and dead zombies (denoted R for *removed*[2]). A zombie can only bite its neighbours that are in state S; conversely, a survivor can only become a zombie if one or more of her neighbours is a zombie.

Let us now add numbers to this description. During an infinitesimal time period $[t, t + dt)$, a survivor with one Z-neighbour has a probability $\alpha{\cdot}dt$ to be bitten by the latter and to become a zombie. For a survivor with n Z-neighbours, the probability would be $n{\cdot}\alpha{\cdot}dt$.

Survivors also have their say. During the same infinitesimal time period $[t, t + dt)$, a zombie with m S-neighbours has a probability $m{\cdot}\beta{\cdot}dt$ to be definitively[3] killed by one of them, henceforth being indefinitely confined to the R state.

We now have everything in hand to model a simple zombie invasion. More complexity can be considered and Section 4 will explore some realistic additions to our basic model. For now, we will limit ourselves to these simple rules.

[2] Or rotting in peace.

[3] Without any possible comeback to “life.”

1.4 Monte Carlo simulations

Now that the rules of the game are set, we can write a computer program to diligently use them. This program has three main tasks to perform: build a network, set the initial conditions and apply the zombie propagation rules.

We have already seen in Section 1.2 how to build a network of known size N from its degree distribution $\{p_k\}$. Once the network is built, the initial state of each of its vertices must be chosen. We will specify this initial condition through the proportion ϵ of the population that starts out as zombies. In other words, all vertices are initially survivors except for ϵN randomly selected ones that are zombies. There are initially no removed individuals.

In order to apply the rules themselves, we discretize time into small intervals of length $\Delta t > 0$. Although this quantity is finite (not infinitesimal), we use it instead of dt in the probabilities of Section 1.3. If Δt is sufficiently small, the resulting dynamics should be a good approximation of the continuous time dynamics[4].

During each of these time intervals, we count how many Z-neighbours each of the survivors has. There is then a probability $n \cdot \alpha \cdot \Delta t$ for each survivors with n Z-neighbours to become a zombie at the next time interval. Similarly, for each zombie, we count the number m of its S-neighbours. The zombie will then be killed (sent to the R state) at the next time interval[5] with probability $m \cdot \alpha \cdot \Delta t$. This is repeated for as many time intervals as required, i.e., until there are no remaining edges linking S and Z vertices together.

We now have a complete procedure (program) providing the state of each vertex at any time t. Since the network construction, the assignment of initial conditions and the propagation rules are all probabilistic in nature, two different realizations of the process (executions of the program) will typically lead to different results. A process of this kind, relying on randomness, is called a *Monte Carlo simulation*.

Usually, one wants to perform many different Monte Carlo simulations using the same parameters in order to obtain reliable statistics about the model's predictions, such as the mean number of zombies at a given time. How many simulations are required? Depending on the problem and the precision required, the answer may range from a few hundred to billions or more.

While numerical simulations are a rather easy way to obtain results, they also have their disadvantages. For instance, it may happen that the total number of simulations to be performed becomes prohibitively high or that the average length of each simulation becomes too long. But perhaps the greatest flaw of such a "brute force" approach is the lack of insight that we gain from it. We do get results, but they do not offer a good grasp on the underlying mechanics. To address this issue, we will use in the next section a completely different approach that provides better insights on the dynamics of the invasion.

2 Mathematical zombies: a theory of the undead

In this section, we basically want to do the same thing as in Section 1.4, but without its principal drawbacks. The alternative approach is based upon rewriting the problem in terms of a set of ordinary differential equations (ODEs). Some approximations need to be done, but the results are in perfect (well, almost) agreements with those obtained by the more direct approach of Section 1.4. As a bonus, however, what is actually happening in the system comes to light (no pun intended).

[4] Note that there are better ways to apply these kind of rules, but this will be sufficient at present.

[5] All changes are applied simultaneously before the next time step.

2.1 The variables of the system

We shall now write an ensemble of equations following the propagation of a zombie outbreak on a social network. One may first be inclined to simply follow the behaviour of the fraction of the population at time t who have survived, $S(t)$, of those who are now zombies, $Z(t)$, and of those who are now left rotting, $R(t)$. However, this approach pretty much nullifies the original incentive for using a network structure in the first place; i.e., to consider the importance of heterogeneous human behaviour on the evolution of the population.

Logically, we need at least to follow the behaviour of survivors $S_k(t)$ and zombies $Z_k(t)$ for each degree k. Since we define our networks through their degree distribution, throwing degrees away would be a waste. However, this is still far from perfect. Consider how different the behaviour of an individual hidden with ten friends would be from the behaviour of an individual surrounded by ten zombies! We therefore need to go a step further and differentiate survivors and zombies according to both the state and the number of their neighbours. To that effect, let $S_{m,n}(t)$ and $Z_{m,n}(t)$ be the proportion of individuals at time t who are survivors or zombies and are currently in contact with m survivors and n zombies.

Note that the total fraction of survivors and zombies at time t may easily be obtained from these quantities by summing over both m and n:

$$S(t) = \sum_{m,n} S_{m,n}(t) \tag{1}$$

$$Z(t) = \sum_{m,n} Z_{m,n}(t)\,. \tag{2}$$

Moreover, the fraction of removed can also be obtained from the fact that $S(t)$, $Z(t)$ and $R(t)$ are fractions of individuals; as such, they should all sum to 1:

$$R(t) = 1 - S(t) - Z(t)\,. \tag{3}$$

Hence, the knowledge of $S_{m,n}(t)$ and $Z_{m,n}(t)$ for all times t solves our model.

2.2 Moment-closure approximation

An astute reader might have noticed a slight problem from the fact that we do *not* know everything about the network. The following example illustrates the essence of this problem.

You are currently one of the $S_{1,0}$: a survivor with one S-neighbour and no Z-neighbours. You are therefore safe — or are you? What if your neighbour becomes a zombie? If he has some Z-neighbours, this may very well occur! On the other hand, if your neighbour is a $S_{1,0}$ like you, you form an isolated pair protected from the zombie invasion. Paranoia and suspicion are key factors when it comes to surviving a zombie apocalypse, and you are thus entitled to know who, apart from you, is connected to your friends.

In Section 2.1, we could very well have chosen a more complicated set of variables to explicitly take these situations into account. However, no matter how much you know about the network, you could always know more: who are the neighbours of the neighbours of your neighbours? And what about their neighbours?

Increasing complexity has a cost, one that may ruin our ability to obtain anything useful. Arguably, a wise choice is to track only the neighbours of each vertex and then *infer* the neighbourhood of these neighbours from the total available information.

Considering our previous example, let us try to guess the state of your unique survivor friend. Once again, when you do not know much, the best choice is to use all the information you have and nothing else. Should we say that he is in the state $S_{m,n}$ with m and n chosen at random? No, because we know more than that. For example, if we know that $S_{5,2}(t) = 0$, your neighbour is certainly not[6] in the state $S_{5,2}$ since nobody is in that state right now. The probability for your neighbour to be in the state $S_{m,n}$ at time t is thus proportional to the population in the state $S_{m,n}(t)$.

Moreover, your neighbour cannot have $m = 0$, i.e., no S-neighbour, since you are there. In fact, the probability for your neighbour to be in the state $S_{m,n}$ at time t is proportional to m: the more S-neighbours he has, the more likely you are one of them[7].

This chain of reasoning leads to the fact that the probability for an individual to be in the state $S_{m,n}$ at time t knowing that he is a survivor and has at least one S-neighbour must be proportional to $m \cdot S_{m,n}(t)$. Requiring normalization, this probability is then

$$\frac{mS_{m,n}(t)}{\sum_{m',n'} m' S_{m',n'}(t)} . \tag{4}$$

It is now quite easy to obtain quantities such as the mean number of Z-neighbours that each S-neighbour of a survivor has:

$$\langle z(t) \rangle_{s,s} = \frac{\sum_{m,n} n \cdot mS_{m,n}(t)}{\sum_{m',n'} m' S_{m',n'}(t)} . \tag{5}$$

This is called a *moment-closure approximation*: basically, the idea is to guess the higher moments of a distribution such that they are consistent with the already-known lower-order moments.

Other moments may be obtained the same way. Noting $\langle i(t) \rangle_{j,k}$ the mean number of neighbours in state i of a vertex in state j, itself being a neighbour of a vertex in state k, we have

$$\langle s(t) \rangle_{z,s} = \frac{\sum_{m,n} m \cdot mZ_{m,n}(t)}{\sum_{m',n'} m' Z_{m',n'}(t)}$$

$$\langle s(t) \rangle_{z,z} = \frac{\sum_{m,n} m \cdot nZ_{m,n}(t)}{\sum_{m',n'} n' Z_{m',n'}(t)}$$

$$\langle z(t) \rangle_{s,z} = \frac{\sum_{m,n} n \cdot nS_{m,n}(t)}{\sum_{m',n'} n' S_{m',n'}(t)} .$$

The problem raised at the beginning of this section is now solved. We still do not know everything about the network, but with these quantities at hands, we have enough to write differential equations governing the evolution of $S_{m,n}(t)$ and $Z_{m,n}(t)$.

[6] Well, "*almost surely not*" would be more accurate, but this is another story.

[7] This is why, on average, your friends always have more friends than you do.

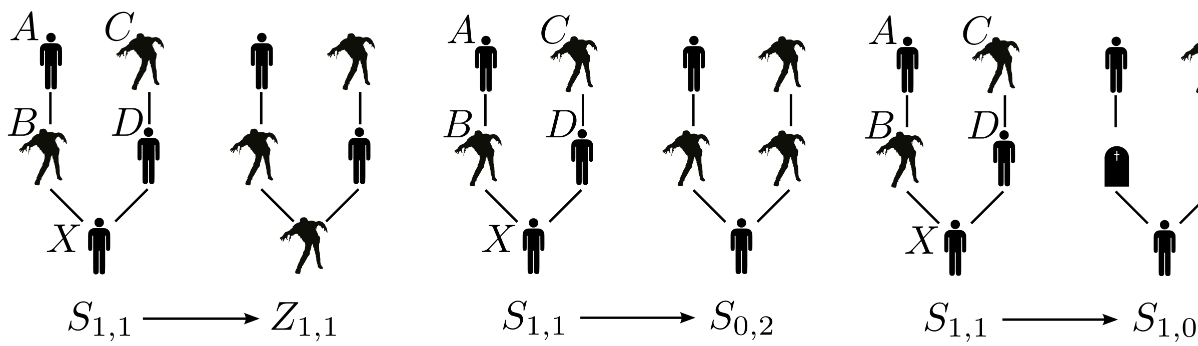

(a) If B bites X... (b) If C bites D... (c) If X or A kill B...

Figure 2: Processes affecting the population of survivors. Transitions are shown for individual X.

2.3 Writing the dynamical system

We shall now consider every possible processes through which an individual might go from one state to another, thereby changing the population of its initial and new states. We then write ODEs providing the rate of change of $S_{m,n}$ and $Z_{m,n}$ for the time-dependent susceptible and zombie nodes[8].

Here are the transitions that can affect a vertex X together with their occurence probability during a time interval $[t, t + dt)$.

(i). *X, a survivor, is bitten by a zombie.* As seen in Section 1.3, this event causes X to become a zombie. Hence, the transition $S_{m,n} \to Z_{m,n}$ occurs with a probability $n\alpha S_{m,n}dt$. An example case is shown in Figure 2a.

(ii). *A surviving neighbour of X is bitten.* This is a side effect of the event (i). Since the state in which X resides depends on the state of its neighbours, the state of X has to change when one of its S-neighbours becomes a zombie. We thus have the transition $S_{m,n} \to S_{m-1,n+1}$ with probability $m\alpha\langle z\rangle_{s,s}S_{m,n}dt$ if X is a survivor and $Z_{m,n} \to Z_{m-1,n+1}$ with probability $m\alpha\langle z\rangle_{s,z}Z_{m,n}dt$ if X is a zombie. Notice the use of $\langle z\rangle_{s,s}$ and $\langle z\rangle_{s,z}$ in order to know the average number of Z-neighbours the S-neighbours of X have. An example is shown in Figure 2b.

(iii). *X, a zombie, is killed by a survivor.* This causes X to become removed through the transition $Z_{m,n} \to R$ with probability $m\beta Z_{m,n}dt$. Since we do not explicitly track removed individuals [but instead obtain them through equation (3)], we simply decrease the population of $Z_{m,n}$ (Figure 2c).

(iv). *A zombie neighbour of X is killed.* This is a side effect of event (iii), not unlike (ii) was to (i). We have the transition $S_{m,n} \to S_{m,n-1}$ with probability $n\beta\langle s\rangle_{z,s}S_{m,n}dt$ if X is a survivor and $Z_{m,n} \to Z_{m,n-1}$ with probability $n\beta\langle s\rangle_{z,z}Z_{m,n}dt$ if X is a zombie. An example is shown in Figure 2c.

[8] From here on, we will lighten the notation by removing the time dependence "(t)" on quantities that clearly vary in time.

The net effect of these transitions on $S_{m,n}$ and $Z_{m,n}$ gives rise to a system of ODEs of the form

$$\begin{aligned} \frac{d}{dt}S_{m,n} &= \langle z\rangle_{s,s}\cdot\alpha\,[(m+1)S_{m+1,n-1}-mS_{m,n}] \\ &\quad+\langle s\rangle_{z,s}\cdot\beta\,[(n+1)S_{m,n+1}-nS_{m,n}]-n\alpha S_{m,n} \end{aligned} \tag{6}$$

$$\begin{aligned} \frac{d}{dt}Z_{m,n} &= \langle z\rangle_{s,z}\cdot\alpha\,[(m+1)Z_{m+1,n-1}-mZ_{m,n}] \\ &\quad+\langle s\rangle_{z,z}\cdot\beta\,[(n+1)Z_{m,n+1}-nZ_{m,n}]+\alpha nS_{m,n}-m\beta Z_{m,n}\,. \end{aligned} \tag{7}$$

The degree distribution $\{p_k\}$ and the initial fraction of zombies ϵ are introduced into the system through the initial conditions

$$S_{m,n}(0) = (1-\epsilon)\,p_{m+n}\binom{m+n}{n}\epsilon^n(1-\epsilon)^m \tag{8}$$

$$Z_{m,n}(0) = \epsilon\,p_{m+n}\binom{m+n}{n}\epsilon^n(1-\epsilon)^m\,. \tag{9}$$

The binomial coefficients come from the fact that a vertex of degree k has probability $\binom{k}{n}\epsilon^n(1-\epsilon)^{k-n}$ to have n Z-neighbours.

Using these initial conditions, we can solve (i.e., by means of numerical integration[9]) the system of ODEs given by equations (6) and (7) to trace the future of a zombie outbreak on our social network.

3 Results: does humanity have the slightest chance?

We now compare the approaches we obtained in Sections 1.4 and 2.3 before extracting meaning from them. But most importantly, this is where we see if humanity has a chance against a zombie invasion.

3.1 Choosing the victims

Before proceeding, we must decide the size of the population to model. A small city of $N = 10\,000$ people should suffice, and let us say that one percent of them are already zombies ($\epsilon = 0.01$), none of which have yet been killed.

What is the degree distribution $\{p_k\}$ in that city? If we say that everybody has equal chance to be connected with everybody and, in addition, we know λ, the mean number of acquaintances people have, then $\{p_k\}$ is provided by the Poisson distribution for large populations (Figure 3):

$$p_k = \frac{\lambda^k e^{-\lambda}}{k!}\,. \tag{10}$$

However, real human populations behave very differently: some people have much more neighbours than the average. There are not many of them, but they disproportionately influence any spreading processes occuring on the network. The *power-law* distribution, with parameter $\tau > 0$ and

[9] For those not familiar with numerical integration techniques, see the introductory notes presented in the appendix of this chapter.

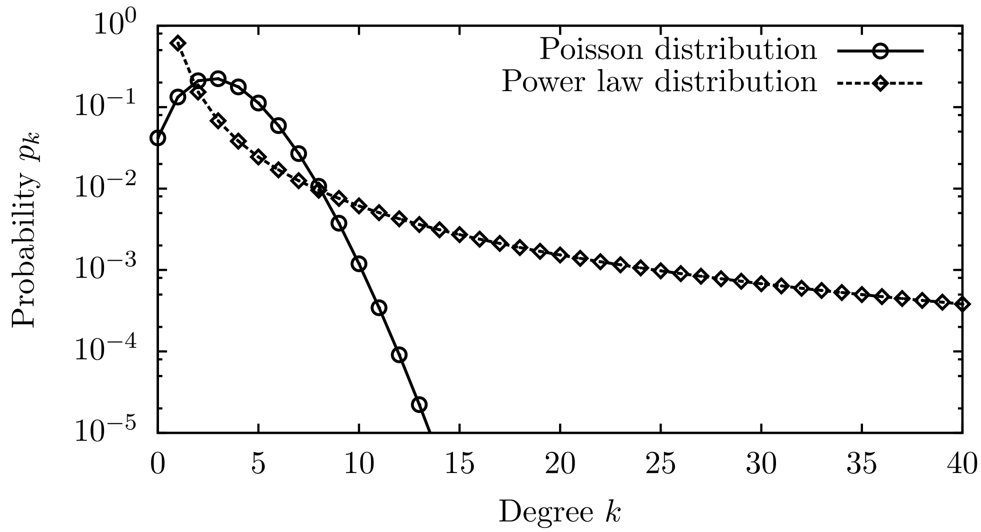

Figure 3: Two degree distributions are considered on this semi-log plot. Although both share the same average degree $\lambda = 3.17$, the power-law distribution ($\tau = 2$, $k_{\max} = 100$) falls very much slower than the Poisson distribution. The number of friends each individual has in the Poisson network does not vary much from one person to another. In the power-law distributed network, some people *really* have a lot of friends.

truncated at $k_{\max}$, namely

$$p_k = \begin{cases} 0 & \text{if } k = 0 \\ \dfrac{k^{-\tau}}{\sum_{k'=1}^{k_{\max}} (k')^{-\tau}} & \text{if } 1 \leq k \leq k_{\max} \\ 0 & \text{if } k > k_{\max} \end{cases} \tag{11}$$

exemplifies this behaviour (Figure 3).

We will use these two distributions and collect the corresponding results. However, to make a fair comparison (to compare "apples with apples" some might say), we use $\lambda = 3.17$, $\tau = 2$ and $k_{\max} = 100$ such that both distributions have the same average degree λ.

3.2 Body count

Our two approaches, i.e., Monte Carlo simulations and integration of our system of ODEs, are applied to the two populations differing by their degree distribution. Figure 4 shows[10] typical results for the parameters $\alpha = 5$ and $\beta = 1$. These parameters represent a (sadly likely) case where zombies are much better at biting survivors than the latters are at killing zombies.

The first observation is the *striking* agreement between Monte Carlo simulations (symbols) and the integrated system of ODEs (curves) for both populations [4a and 4c]. What are the implications of such a match? When a sufficiently small time increment Δt is used for Monte Carlo simulations, the results should be arbitrarily close to an exact realization of the rules we chose to model (i.e.,

[10] The networks snapshots were produced using Eytan Adar, *GUESS: The Graph Exploration System*, `http://graphexploration.cond.org`.

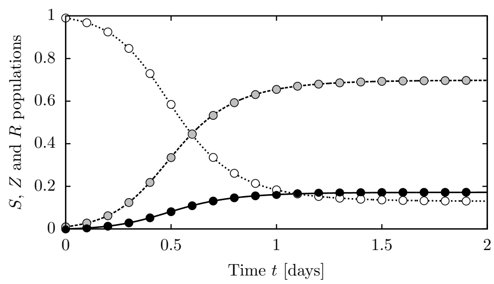

(a) Evolution for Poisson distribution.

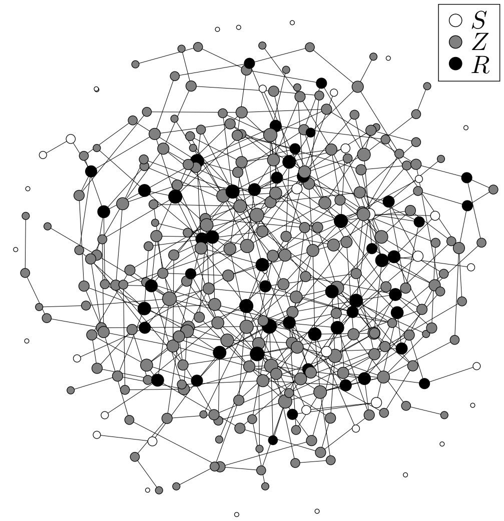

(b) Final state for Poisson.

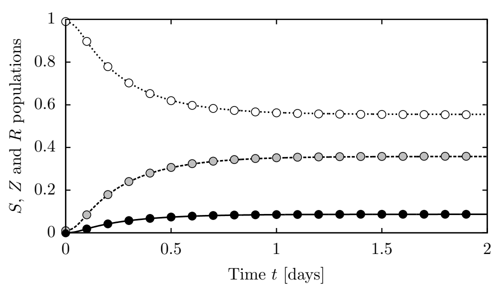

(c) Evolution for power-law distribution.

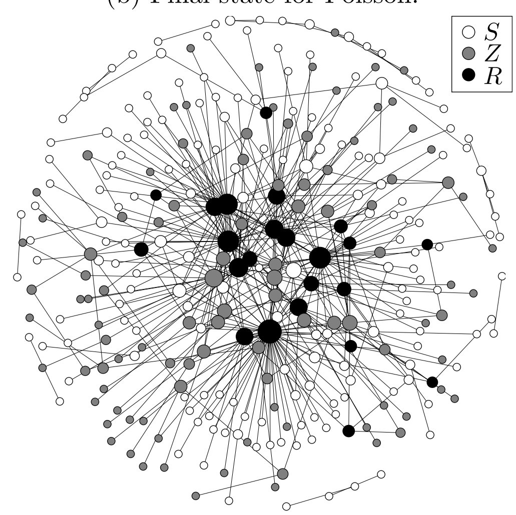

(d) Final state for power-law.

Figure 4: Time evolution of the zombie invasion for networks with (a) Poisson and (c) power-law degree distributions (same average degree $\lambda = 3.17$). Curves (dotted S, dashed Z and solid R) show results for the ordinary differential equation formalism while symbols (clear S, grey Z and black R) represent those obtained through averaging over 100 Monte Carlo simulations. Figures (b) and (d) show the final state of a single Monte Carlo simulation performed in the same conditions as (a) and (c) respectively, except that a smaller network size (300 vertices) was used in order to improve visibility. The visual representations are produced by the GUESS software using an algorithm that distinguishes vertices by their state (colour) and their degree (size).

both the network structure and the propagation rules). Thus, if the integrated system of ODEs agrees with the simulations then they in turn accurately reflect the chosen rules.

This accuracy, despite simplifications and approximations while building the system of ODEs, means that the ignored information does not have much impact on the propagation dynamics of the whole system after all. We have therefore gained the important insight that everything of importance has been included in the model.

A second observation further refines this insight: there is significant differences between the results for the two degree distributions. In a non-network model, these choices would both give the same results since they lead to identical average degree for vertices. This sole fact justifies by itself the need for a network-based approach to the problem.

A third important observation is that both scenarios explored in Figure 4 eventually reach an equilibrium where survivors and zombies *coexist*. This may be surprising to some readers who may be too closed-minded to consider cohabitation with the undead, but it is an unavoidable consequence of the set of rules we chose and has previously been observed in some particular cases[11,12].

No equilibrium can exist in this model as long as there are survivors with zombie neighbours: either the survivor would eventually kill the zombie or be bitten by the latter. Hence, survivors can coexist with zombies in the equilibrium state only if the neighbours of their neighbours of their neighbours... well, all of them, are survivors *or* if the only paths leading to zombies have to go through removed vertices. Hence, sub-networks of survivors and sub-networks of zombies may coexist in the same network as long as they are separated by a "no man's land" of (probably) rotting and (really) dead corpses. In order to get rid of zombies, human populations should therefore develop very efficient tools to eliminate them at the onset of the invasion, or otherwise should bring themselves to consider and to accept such untypical cohabitation.

An observation that illustrates the uniqueness of a zombie invasion is that things are actually *improving* (i.e., there are more survivors at the end) with a power-law degree distribution compared to the Poisson degree distribution. The presence of high-degree vertices, often called *hubs*, usually makes things worse for "normal" diseases. Consider for example the impact of sex workers on sexually transmitted infections or of hospitals/schools for pulmonary infections. Once infected, these hubs quickly dispatch the disease to a lot of people, some of these also being hubs, and this results in many individuals getting infected.

Things are different for zombies, and the difference lies in the very fact that survivors may obliterate zombies. Once a hub is zombified, it can surely bite its many neighbours, but these neighbours also get a chance to eradicate the zombie hub. Most of the big vertices, hubs, are thus eliminated by their neighbours [see Figure 4d]. These removals effectively reduce the average degree of the remaining vertices, since their degree was originally large compared to the mean. In a more "balanced" situation such as the Poisson distribution [Figure 4b], hubs are mostly absent. When a zombie dies, it was just an average zombie and may very well be replaced by another. Had we chosen $\beta = 0$ (or a negligibly small value), the usual effect of hubs would have been observed.

4 The social zombie: adapting for realism

There are some fundamental problems with the dynamics we chose to model. For example, survivors will fight with any number of zombie neighbours, until the undead perish or until the survivors themselves join the nightmarish hordes. In real life, an individual will most likely choose to take flight when the situation becomes too dire and survivors will then try to regroup while avoiding

[11] Shawn of the Dead, Edgar Wright, 2004
[12] Land of the Dead, George Romero, 2005

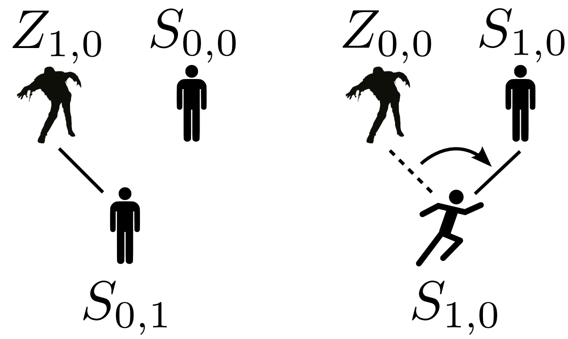

Figure 5: A case of flight. A survivor runs away from a zombie (no longer being its neighbour) and reconnects to a random new neighbour, here another survivor.

zombies. On the other hand, the undead are also usually seen in hordes. Zombies are always inclined to give a hand to their fellow undead friends and help them hunt their remaining targets. This cooperative behaviour between zombies will be used to emulate the typical formation of zombie hordes.

While the networks we used up to now were *static*, i.e., the same edges were always linking the same vertices, the introduction of new rules will allow the structure of the network itself to change. By using *adaptive networks*, we are moving one step further away from the homogeneous approximation. This level of details comes at additional cost, but as it will be shown, it is well worth the spending.

4.1 The new rules of the game

The rules introduced in Section 1.3 have to be adapted. Survivors are now allowed to run away (flight) when facing a zombie and to form groups into barricades for a better defence against the threat. On the other hand, zombies may form hordes improving their collective biting abilities.

During an infinitesimal time interval $[t, t+dt)$, a survivor facing a zombie has a probability γdt to flee from it and to stop at the next individual she encounters (see Figure 5). The survivor then disconnects from the zombie and reconnects to either a survivor ($S_{m,n} \rightarrow S_{m+1,n-1}$) or to a zombie ($S_{m,n} \rightarrow S_{m,n}$). Like in Section 2.3, this causes "side effect transitions" both to the zombie he flees from and to its new neighbour. The new neighbour will be in a state whose probability is directly proportional to the population of that very state.

Both rules for survivors killing zombies or zombies biting survivors are adapted such that an individual is better at attacking the enemy when she is helped by neighbours of her own type (S or Z). During the time interval $[t, t+dt)$, a zombie with n fellow flesh-eaters can bite a survivor with probability $\alpha(1 + h \cdot n)dt$, just as a survivor defending herself alongside m friends will kill a zombie with probability $\beta(1 + h \cdot m)dt$. We will refer to h as the *H-factor*[13], a parameter determining how useful friends are in combat. Again, side effect transitions are affected likewise.

The derivation of the system of ODEs is left as an exercise. Here we provide only the final

[13] "H" may stand for *hordes*, *help*, *hunting*, *hellbent* or even *Hades* itself. Pick your poison.

result:

$$\begin{aligned}\frac{d}{dt}S_{m,n} &= (h\langle z\rangle_{z,s}+1)\,\langle z\rangle_{s,s}\alpha\,[(m+1)S_{m+1,n-1}-mS_{m,n}] \\ &\quad+(h\langle s\rangle_{s,z}+1)\,(\langle s\rangle_{s,z}-1)\,\beta\,[(n+1)S_{m,n+1}-nS_{m,n}] \\ &\quad+(1+hm)\,\beta\,[(n+1)S_{m,n+1}-nS_{m,n}]-(h\langle z\rangle_{z,s}+1)\,n\alpha S_{m,n} \\ &\quad+\gamma\frac{S}{S+Z}\,[(n+1)S_{m-1,n+1}-nS_{m,n}]+\gamma\langle z\rangle_s\frac{S_{m-1,n}-S_{m,n}}{S+Z}\end{aligned} \tag{12}$$

$$\begin{aligned}\frac{d}{dt}Z_{m,n} &= (h\langle z\rangle_{z,s}+1)\,(\langle z\rangle_{s,z}-1)\,\alpha\,[(m+1)Z_{m+1,n-1}-mZ_{m,n}] \\ &\quad+[(n-1)h+1]\,(m+1)\alpha Z_{m+1,n-1}-(hn+1)\,m\alpha Z_{m,n} \\ &\quad+(h\langle s\rangle_{s,z}+1)\,\langle s\rangle_{z,z}\beta\,[(n+1)Z_{m,n+1}-nZ_{m,n}] \\ &\quad-(h\langle s\rangle_{s,z}+1)\,m\beta Z_{m,n}+(h\langle z\rangle_{z,s}+1)\,n\alpha S_{m,n} \\ &\quad+\gamma\,[(m+1)Z_{m+1,n}-mZ_{m,n}]+\gamma\langle z\rangle_s\frac{Z_{m-1,n}-Z_{m,n}}{S+Z}\,.\end{aligned} \tag{13}$$

For convenience, we defined the average number of Z-neighbours a survivor has:

$$\langle z\rangle_s=\sum_{m,n} nS_{m,n} \tag{14}$$

which appears in the equations since the probability for any individual to encounter a survivor running away from a zombie depends on that quantity.

Note that the new system encompasses the previous model since equations (12) and (13) reduce to equations (6) and (7) when one sets the flight rate equal to zero ($\gamma = 0$) and chooses a null *H-factor* ($h = 0$).

4.2 Hope is the last thing to die, isn't it?

Figure 6 is obtained by performing Monte Carlo simulations and integrating the system of ODEs corresponding to the new rules introduced in Section 4.1. From a quick comparison of the three different cases studied, it is clear that the *H-factor* (i.e., the cooperation effect) plays out to the advantage of the zombies. We could expect it to be the other way around if survivors were more efficient at killing zombies than the latter are at biting (if only...). On the other hand, allowing survivors to flee (i.e., $\gamma > 0$) from the zombies gives them a better life expectancy, but only postpones the inevitable. In fact, all three scenarios eventually lead to an equilibrium where around 15% of the population manages to survive by barricading themselves against the undead hordes. Again, our results suggest similar conclusions: human populations should develop efficient weapons and fighting tactics in order to eradicate these nightmarish invaders before they take over the world.

Notice that the fit between the Monte Carlo simulations and the analytical predictions is somewhat less satisfying in the scenario where flight is possible. Great insights can be gained from this simple observation: something we either approximated or neglected when obtaining the system of ODEs becomes important when flight is allowed.

The culprit is most likely correlation. In our system of ODEs, when a survivor flees from a zombie and then encounters vertex X, we add a new S-neighbour to vertex X. Later on, we consider this survivor as the *average survivor*. This is where the shoe pinches: he is not the average survivor, but the average survivor who just ran away from a zombie, and this additional information tells us something about his state. He would probably not have been able to get this far had he been

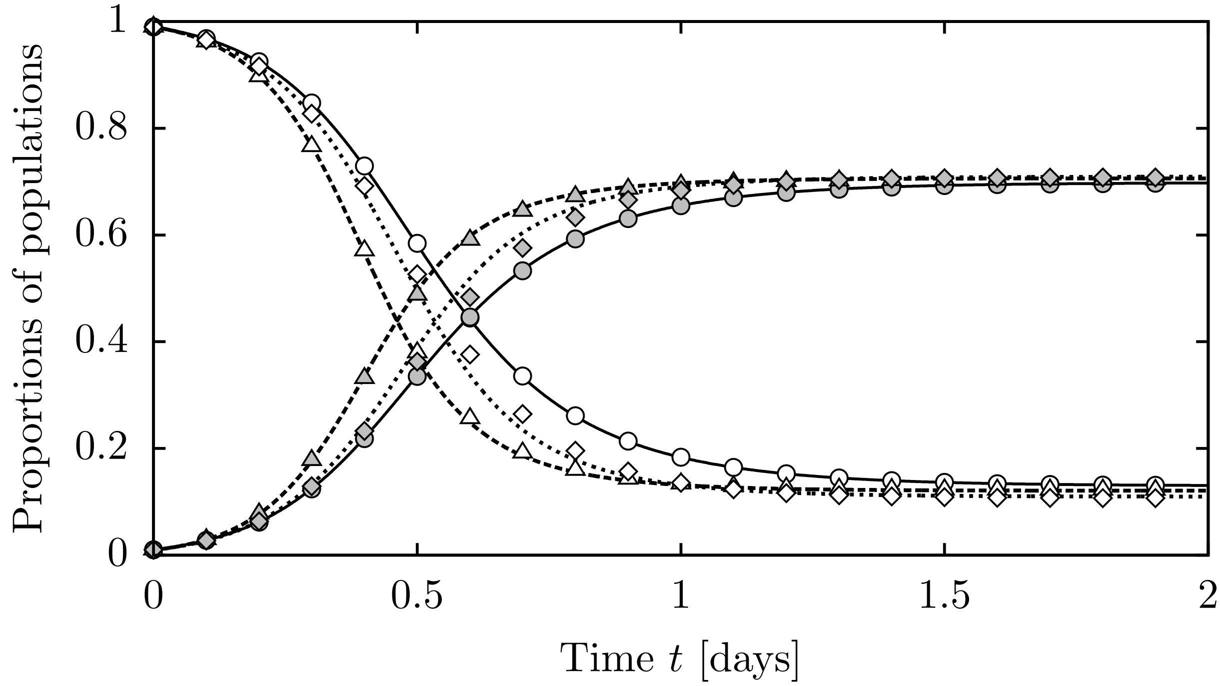

Figure 6: Monte Carlo (symbols) and integration of the system of ODEs [equations (12) and (13)] (curves) results for the new set of rules. The "old" parameters are the same as those of Figure 4a, and the Poisson degree distribution was used. For new parameters, three different cases are studied: $h = 0$ and $\gamma = 0$ (solid curves and ○); $h = \lambda^{-1}$ and $\gamma = 0$ (dashed curves and △); and $h = \lambda^{-1}$ and $\gamma = 2$ (dotted curves and ◇). Survivors are shown in white and zombies in grey.

in contact with 100 zombies, nor would he be fleeing had he not been in contact with at least one zombie in the first place. The fact that our equations do not consider these correlations is probably the principal cause of the observed differences with the Monte Carlo simulations.

In order to improve the results, one could embark on the tedious endeavour of including the state of the neighbours of the neighbours of each vertex in the description of the system. Doing so would require something like

$$S_{\substack{m_{0,0}, m_{0,1},\, m_{0,2},\, ...,\, m_{1,0},\, m_{1,1},\, ... \\ n_{0,0},\, n_{0,1},\, n_{0,2},\, ...,\, n_{1,0},\, n_{1,1},\, ...}} \quad \text{and} \quad Z_{\substack{m_{0,0}, m_{0,1},\, m_{0,2},\, ...,\, m_{1,0},\, m_{1,1},\, ... \\ n_{0,0},\, n_{0,1},\, n_{0,2},\, ...,\, n_{1,0},\, n_{1,1},\, ...}} \tag{15}$$

where each $m_{i,j}$ represent the number of survivor neighbours who themselves have i survivor and j zombie neighbours ($n_{i,j}$ playing the same role for zombie neighbours).

Since a differential equation would be required for each of these quantities, it is easy to imagine how quickly the complexity of this sort of models can explode. However, considering that the error observed in the results was not that critical in the first place, the simpler approach that we have used was appropriate to the problem we chose to tackle.

5 A conclusion on networks...

This was only the tip of the iceberg. As we have witnessed in this chapter, network modelling (namely, the inclusion of a topology in a model of interaction between a large number of elements) is useful when considering the effects of local dynamics on the global evolution of the system. What does this mean for zombie modelling? The model developed here allowed us to consider multiple local effects:

Heterogeneity of behaviours... and fates. Some individuals can be (or get) disconnected from the rest of the world while others might end up in the middle of the zombie apocalypse.

Local environment. Individuals are not affected by the global state of the world, but only by their immediate environment (this can allow a single zombie to start a very virulent, but localized outbreak).

Social behaviours. Networks do not have to be static: humans can flee from the zombies and regroup, while the undead will hunt in hordes.

However, a multitude of other similar effects could have been (and should be) incorporated in a more complex model:

Types of individuals. Not all humans will react the same in a zombie invasion (e.g. the behaviour of a soldier is likely to differ from that of a child) and network modelling is perfectly suitable for considering such heterogeneity.

Social structure. We have given very basic behaviour to both the living and the dead, but the formation of more complex social structure could be taken into account.

Note that, while network modelling is a natural approach to the description of individual and heterogeneous behaviours, the inclusion of structure in the problem usually complicates its treatment. Furthermore, complexity usually increases rapidly in network theory. This means that our little exercise in network modelling would have been an entirely different story had we considered an even more complex system.

...and zombies

What have we learned from this exercise that could help us survive the imminent zombie apocalypse?

Firstly, we have learned that, in a scenario where the set of rules that we have chosen applies, any invasion must be stopped at its very beginning or a significant fraction of the population will be either dead or zombified, thus marking the end of the supremacy of humanity on Earth. We therefore stress once again the importance of developing efficient and powerful anti-zombie weapons and defenses[14].

Secondly, including a cooperation effect in our model has been shown to *accelerate* the zombie invasion (if $\alpha > \beta$) while our results show that flight only slows down the progression of the invasion without significantly modifying its outcome. Therefore, if we consider fighting back the undead in a "human-to-zombie" fashion, we should use state of the art fighting techniques in order to bring the cooperation effect to our advantage (i.e., $\beta > \alpha$)[15].

Finally, within the context in which our model has been developed, it has been shown that any invasion can be stopped, but cannot be reversed. Once a succesful barricade is established, we can expect to find a certain number of roaming, isolated zombies. To wipe out these remaining threats, we believe it is necessary to consider more than fighting at the individual level. Optimistically, efforts should be deployed in developing antidotes and vaccines. On the other hand, the use of weapons of mass destruction should be considered once a "stable" condition is reached. In light of the results presented in this chapter, such strategies may well be our only chance to take these poor souls back to Hell from whence they spawned... until next time.

> *Yea, though I walk through the valley of the shadow of death, I will fear no evil: for thou art with me, you mathematical-physicist; thy model and thy simulation they comfort me*[16].

[14] In our model, such equipment would result in reducing α and increasing β.

[15] Again, in our model, such a strategy could be studied by using different values of the *H-factor* for survivors and zombies.

[16] Any ressemblance with Psalm 23:4 of the Original King James Bible is purely coincidental.

Appendix: Numerical Integration

In this chapter, we have followed a modelization approach based on a large set of nonlinear ordinary differential equations (ODEs) to describe the dynamics of a zombie outbreak on a social network. Solving this system of ODEs by analytical means would be impossible due to its size and complexity. Thus, one has to rely on *numerical integration* to approximate the analytical solution for the time evolution of our system.

In order to illustrate the basics of numerical integration, let us consider a simple one dimensional ODE system:

$$\frac{dx(t)}{dt} = f\big(x(t)\big)\,. \tag{16}$$

with the initial condition $x(t_0)$, and for which we do not know any analytical solution. Consider now the following intuitive reasoning, using an analogy to unidimensional motion. Suppose $x(t)$ is the position of a particle at time t. At this particular time, its speed (in general, the rate of change of the variable) is given by $f\big(x(t)\big)$. Starting from the known initial condition, we are now interested in finding where this particle will be at time $t_0+\delta t$ based on the knowledge of the position $x(t_0)$ and speed $f\big(x(t_0)\big)$ at time t_0. For δt sufficiently small, the speed of the particle will remain approximately constant during this short time step. This results in a linear motion, given by:

$$x\left(t_0 + \delta t\right) = x\left(t_0\right) + \delta t f\big(x(t_0)\big)\,. \tag{17}$$

This procedure can then be used iteratively to generate an approximate solution at later discrete times; i.e.,

$$x\left(t_0 + 2\delta t\right) = x\left(t_0 + \delta t\right) + \delta t f\big(x(t_0 + \delta t)\big) \tag{18}$$

and so on.

This numerical integration scheme is known under the name of *Euler's method* and is schematically illustrated in Figure 7. With this method, the approximate solution has a *local error*, i.e., the error made at each time step δt, of order $\mathcal{O}\left(\delta t^2\right)$. Practically, it means that the numerical approximation gets closer to the real solution as we diminish δt. Formally, this is because equation (17) corresponds to the first two terms of a Taylor series expansion of $x(t_0 + \delta t)$ around t_0:

$$x(t_0 + \delta t) = x(t_0) + \left.\frac{dx(t)}{dt}\right|_{t_0} \delta t + \left.\frac{d^2x(t)}{dt^2}\right|_{t_0} \delta t^2 + \left.\frac{d^3x(t)}{dt^3}\right|_{t_0} \delta t^3 + \mathcal{O}(\delta t^4)\,. \tag{19}$$

Consequently, if we integrate the equation (16) for a total time $T = N\delta t$, the resulting *global error* will be of order $\mathcal{O}(\delta t)$.

Euler's method can also be easily generalized to system of ODEs of more than one dimension, as it was the case in our problem:

$$\vec{x}\left(t_0 + \delta t\right) = \vec{x}\left(t_0\right) + \delta t \vec{f}\big(\vec{x}(t_0)\big)\,, \tag{20}$$

where

$$\vec{x}(t) = \begin{pmatrix} x_1(t) \\ x_2(t) \\ \vdots \\ x_M(t) \end{pmatrix} \quad \text{and} \quad \vec{f}\big(\vec{x}(t)\big) = \begin{pmatrix} f_1\big(x_1(t), x_2(t), \ldots, x_M(t)\big) \\ f_2\big(x_1(t), x_2(t), \ldots, x_M(t)\big) \\ \vdots \\ f_M\big(x_1(t), x_2(t), \ldots, x_M(t)\big) \end{pmatrix}. \tag{21}$$

Other, more accurate, methods are also available. One of the most frequently used family of algorithms for ODE integration is based on the Runge-Kutta method, which uses several points in the interval $[t, t + \delta t]$ to approximate the effective slope of $x(t)$. In this chapter, we have used an even more complicated technique[17] based on a Runge-Kutta algorithm which has an estimated

[17] Sometimes called the *Dormand-Prince method* or, more affectively, *Runge-Kutta 4–5*.

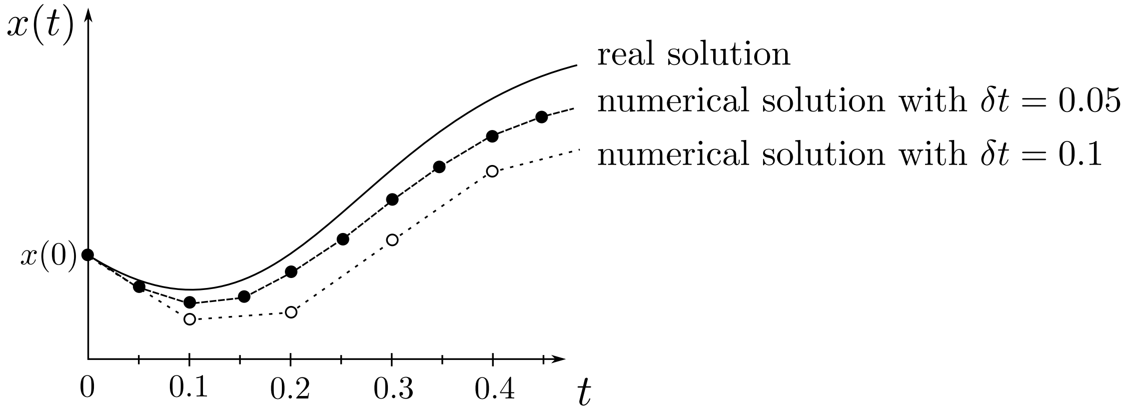

Figure 7: Schematization of Euler's method.

global error of order $\mathcal{O}(\delta t^4)$ and uses an adaptive time step δt in order to keep the error small even when the slope of $x(t)$ is very steep. The interested reader can find a wealth of information in the book *Numerical Recipes* by Press *et al.*

6 Glossary

Adaptive network. A network whose topological structure evolves adaptively with the dynamics taking place on it, thus creating a feedback loop. In other words, the dynamics *on* the network (zombie outbreak) influences the dynamics *of* the network (social contacts), which in turn affect the way the dynamics will further evolve *on* the network.

Degree. The number of neighbours of a given vertex.

Degree distribution. The set of probabilities $\{p_k\}$ for a vertex chosen at random in the graph to have a degree k.

Edge. A line joining two vertices in a graph. In a contact-network model, it represents a social contact between two individuals (e.g. family bonds, friends, etc).

Graph. Mathematical abstract representation of a network, consisting of dots (vertices) and lines (edges).

Moment-closure approximation. A mathematical technique used to approximate higher-order moments of a system with the knowledge of lower-order moments. It is often used to express systems of equations in a closed form.

Monte Carlo simulation. A computer algorithm (simulation) relying on the repeated use of random numbers.

Neighbours. Two vertices linked by an edge.

Network. An interconnected system of *things*. Those things can be people, electrical components, words, computers, etc. The way they are connected together defines the topological structure of the network.

Ordinary differential equation (ODE). A relation involving functions of *only one* independent variable and one or more of their derivatives with respect to that variable.

Static network. a system where the interconnections between elements do not change with time such that its topological structure remains fixed.

Vertex. A dot in a graph. In a contact-network model, it often represents an individual (alive, dead, or in between) in a population.

Further Reading

General literature on networks and network theory

- M.E.J. Newman, A.-L. Barabási, and D.J. Watts, *The Structure and Dynamics of Networks*, Princeton University Press (2006).
- A. Barrat, M. Barthélemy, and A. Vespignani, *Dynamical Processes on Complex Networks*, Cambridge University Press (2008).
- M.E.J. Newman, *Networks: An Introduction*, Oxford University Press (2010).

Contact network epidemiology

- M.E.J. Newman, "Spread of epidemic disease on networks," *Phys. Rev. E* **66**, 016128 (2002).
- M.J. Keeling and K.T.D. Eames, "Networks and epidemic models," *J. R. Soc. Interface* **2**, 295 (2005).
- L.A. Meyers, "Contact network epidemiology: Bond percolation applied to infectious disease prediction and control," *Bull. Am. Math. Soc.* **44**, 63 (2007).
- P.-A. Noël, B. Davoudi, R.C. Brunham, L.J. Dubé, and B. Pourbohloul, "Time evolution of epidemic disease on finite and infinite networks," *Phys. Rev. E* **79**, 026101 (2009).
- A. Allard, P.-A. Noël, L.J. Dubé, and B. Pourbohloul, "Heterogeneous bond percolation on multitype networks with an application to epidemic dynamics," *Phys. Rev. E* **79**, 036113 (2009).
- L. Hébert-Dufresne, P.-A. Noël, V. Marceau, A. Allard, and L.J. Dubé, "Propagation dynamics on networks featuring complex topologies," *Phys. Rev. E* **82**, 036115 (2010).

Adaptive networks

- T. Gross, C.J. Dommar D'Lima, and B. Blasius, "Epidemic dynamics on an adaptive network," *Phys. Rev. Lett.* **96**, 208701 (2006).
- T. Gross and B. Blasius, "Adaptive coevolutionary networks: A review," *J. R. Soc. Interface* **5**, 259 (2008).
- V. Marceau, P.-A. Noël, L. Hébert-Dufresne, A. Allard, and L.J. Dubé, "Adaptive networks: coevolution of disease and topology," *Phys. Rev. E* **82**, 036116 (2010).

Numerical methods

- W.H. Press, S.A. Teukolsky, W.T. Vetterling and B.P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press (2007).