The Social Zombie

Modelling undead outbreaks on social networks

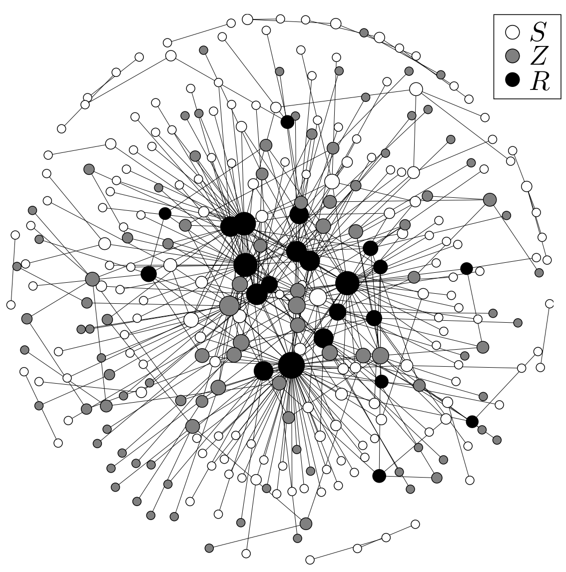

Networks

- We encounter networks all the time:
  - roads
  - public transport
  - internet
  - online social networks
  - facebook
  - satellite networks (GPS)
  - radio (electrical and information networks)
- We can benefit from their advantages and protect ourselves from the disadvantages of networks.

The social network

- People interact with each other for many reasons
  - friendship
  - family
  - sexual partnerships
  - business relations
- At the population level, these interactions sum up to form a giant web...

  ...the social network.

The nature of social interactions

- How individuals are connected to one another in the social network depends on the nature of the interactions
- Within the structure, people may transmit or exchange information, opinions or infectious diseases

  (to name a few)
- The underlying social network shapes the propagation dynamics.

Contact networks

- A contact network is represented by a *graph*
- This is a set of dots linked by lines
- Dots are often called *vertices*
- Lines are called *edges*
- Individuals are represented by vertices
- Two individuals are linked by an edge if they interact with each other.

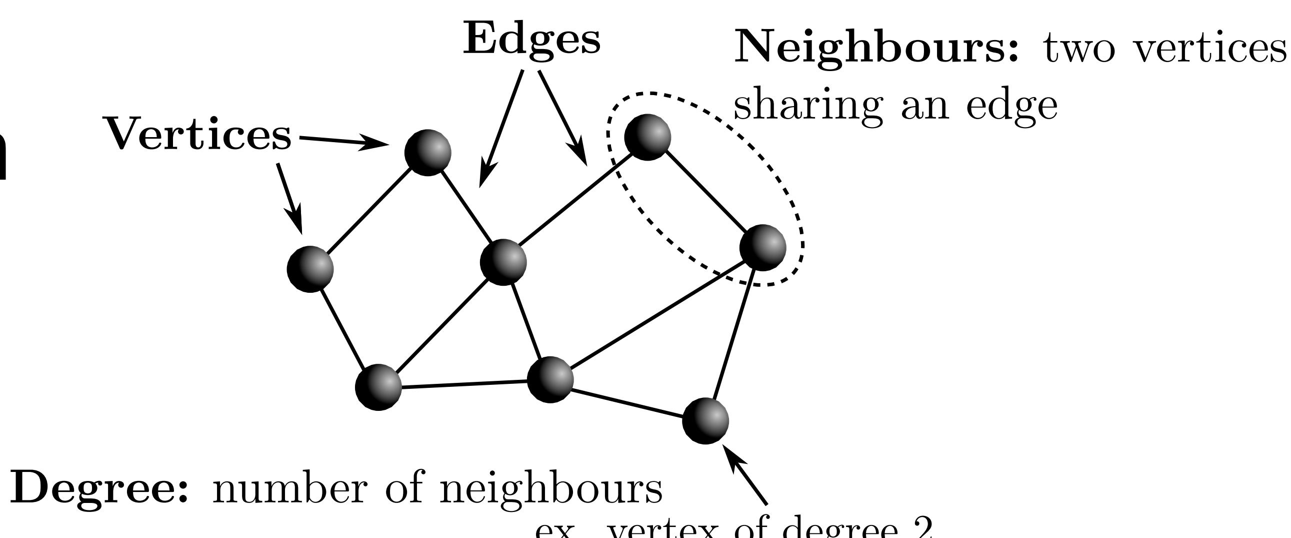

Degree distribution

- Two vertices sharing an edge are called *neighbours*
- The number of neighbours of a given vertex is called its *degree*
- The *degree distribution* of a network is $\{p_k\}=[p_0,p_1,p_2,\ldots]$
- Each p_k gives the fraction of the vertices that have k neighbours
- This is also the probability for a randomly selected vertex to be of degree k.

Determining the degree distribution

- Consider the sample graph from before
- There are
  - no vertices with 0 neighbours
  - no vertices with 1 neighbour
  - 4 vertices with 2 neighbours
  - 2 vertices with 3 neighbours
  - 2 vertices with 4 neighbours
- The degree distribution is $\{p_k\}=[0,0,1/2,1/4,1/4,0,0,...]$.

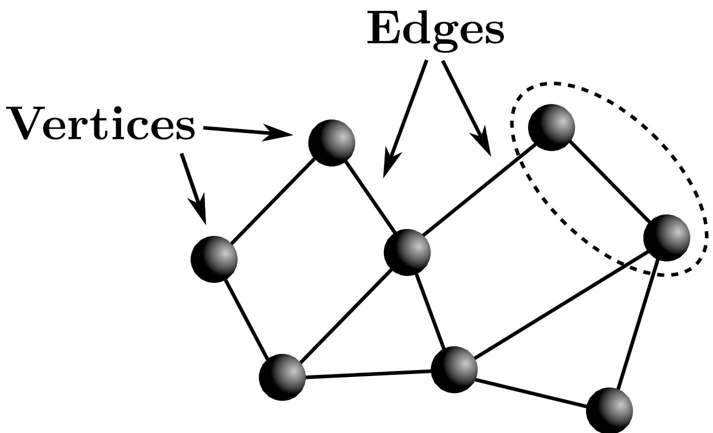

Building a contact network

- The structure of a network contains a lot of information
- To acquire information through a survey, we would have to ask each vertex who its neighbours are
- This might be doable for a very small network, but is usually not possible for larger populations
- Instead, we could survey a fraction of the population and then ask what information this gives us about the whole network.

A partial survey

- Suppose our survey asked 100 people how many neighbours they have

  (ie what is their degree?)
- The only other thing we know is that there are 10,000 people in total
- How do we fill in the blanks?
- ie how do we obtain the degrees of the other 9,900 other vertices?
- And how do we join all the vertices together?

Answering the first question

- Q1: How do we obtain the degrees of the other 9,900 other vertices?
- Easy to answer if we require consistency
- ie if we asked any randomly chosen 100 individuals, they should have the same answers
- In practice, we assume the remaining 9,900 people follow the same probability distribution
- That is, each new vertex is assigned degree k with probability p_k.

p_k: degree distribution

Answering the second question

- Q2: How do we join all the vertices together?
- We could design a complicated process based on
  - friendship
  - encounters on the street
  - who is going to which supermarket
  - etc
- But what if we don't have this information?
- Should we make it up?

Randomly assigning edges

- The best choice is (always) to use the little we do know and nothing that we do not
- In our case, this means forcing vertices to have the degrees we earlier chose
- Then we can randomly assign edges between them
- The edges will have the property that they will satisfy the degree of each vertex, but everything else is random.

An algorithm for randomly assigning edges

- For each of the N vertices, place k pieces of paper in a bag
- Each bears 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 pieces of paper
- Repeat until the bag is empty.

Building a network

- The process of building a network is very similar to the paper-drawing idea
- Of course, everything will be automatic in a computer program
- Thus, we can define networks quite simply
- All we need to know is
  - their size N
  - their degree distribution $\{p_k\}$.

Three kinds of beings

- We can differentiate individuals into three states:
  - susceptibles (S) aka humans
  - infected (Z) aka zombies
  - removed (R) aka dead zombies
- A zombie can only bite its neighbours that are in state S
- Conversely, a human can only become a zombie if one or more of her neighbours is a zombie.

Probabilities in infinitesimal time

- Consider an infinitesimal time period [t,t+dt)
- A survivor with one Z-neighbour has a probability $\alpha \cdot dt$ of being bitten and becoming a zombie
- A survivor with n Z-neighbours has probability $n \cdot \alpha \cdot dt$
- A zombie with m S-neighbours has a probability $m \cdot \beta \cdot dt$ to be definitively killed
- These dead zombies are indefinitely confined to the R state.

S: humans
Z: zombies
R: dead

Initial conditions

- We specify the initial condition through the proportion ϵ of the population that starts out as zombies
- Thus, there are ϵN randomly selected zombies and $(1-\epsilon)N$ humans
- Next, we discretise time into small intervals of length $\Delta t>0$
- This quantity is finite, not infinitesimal like dt
- If Δt is sufficiently small, the resulting dynamics should be a good approximation of the continuous time dynamics.

N: total population

Counting neighbours

- During each Δt time interval, we count how many Z-neighbours each human has
- There is then a probability $n \cdot \alpha \cdot \Delta t$ for each human with n Z-neighbours to become a zombie at the next time interval
- We also count its number m of S-neighbours
- The zombie will be permanently killed at the next time interval with probability $m \cdot \beta \cdot \Delta t$
- This is repeated for as many time intervals as required until there are no S-Z edges.

S: humans
Z: zombies
α: biting probability

Monte Carlo simulation

- We now have a complete procedure to determine the state of each vertex at any time t
- The network construction, the initial condition assignment and the propagation rules are all probabilistic in nature
- Thus, two different realisations of the process will typically lead to different results
- This kind of process, which relies on randomness, is called a *Monte Carlo simulation*.

How many simulations

- We usually want to perform many different Monte Carlo simulations using the same parameters
- This will produce reliable statistics about the model’s predictions
  – eg the mean number of zombies at a given time
- How many simulations are required?
- Depending on the problem and the precision required, it could be a few hundred or more than a billion.

Advantages and disadvantages

- Numerical simulations are an easy way to obtain results
- However, the total number of simulations may become prohibitively high
- Or the average length of each simulation may become too long
- The greatest flaw is the lack of insight that we gain
- Although we do get results, they do not offer a good grasp on the underlying mechanics.

Insight into the mechanics

- To gain insight into the mechanics, we want to write a system of equations
- These will follow the propagation of a zombie outbreak on a social network
- However, we don't want to just follow S(t), Z(t) and R(t) over time, as this would invalidate using a network
- Logically, we need to follow the behaviour of $S_k(t)$ and $Z_k(t)$ for each degree k
- Remember that we define networks through their degree.

S: humans
Z: zombies
R: dead

State and number of neighbours

- However, just knowing the degree isn't enough
- Consider the difference in behaviour of a human hidden with ten friends...

  ...versus a human surrounded by ten zombies!
- Thus, we need to differentiate humans and zombies according to both the state and the number of their neighbours
- Let $S_{m,n}(t)$ and $Z_{m,n}(t)$ be the proportion of humans or zombies in contact with m humans and n zombies.

S: humans
Z: zombies

The total fractions

- The total fraction of humans at time t is

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

- Similarly, the total fraction of zombies at time t is

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

- The fraction of removed can be obtained from the fact that the total sums to 1:

$$R(t) = 1 - S(t) - Z(t)$$

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

S: humans
Z: zombies
R: dead
n: S-neighbours
m: Z-neighbours

We don’t know everything about the network

- Suppose you are a human with one human neighbour and no zombie neighbours
  - ie an $S_{1,0}$
- You appear to be safe, but are you?
- What if your neighbour become a zombie?
- If he has some Z-neighbours, this may occur
- But if your neighbour is also $S_{1,0}$, then you form an isolated pair and are protected
- You are thus entitled to know who, apart from you, is connected to your friends.

S: humans
Z: zombies

Important questions in an apocalypse

- Who are the neighbours of your neighbours?
- What about their neighbours?
- Paranoia and suspicion are key factors when it comes to surviving a zombie invasion!
- However, increasing complexity has a cost
- One that may ruin our ability to obtain anything useful
- A wise choice is to track only the neighbours of each vertex
- Then infer the neighbourhood from the total available information.

Inferring the state of your $S_{1,0}$ neighbour

- Consider the previous example
- Let's try to guess the state of your unique survivor friend
- Is he in the state $S_{m,n}$ with m and n chosen at random?
- No, because we know more than that
- Eg if we know that $S_{5,2}(t)=0$, then your neighbour cannot be in that state
- The probability for your neighbour to be in the state $S_{m,n}$ at time t is proportional to the population in the state $S_{m,n}(t)$.

S: humans
n: S-neighbours
m: Z-neighbours

Chain of reasoning

- Moreover, your neighbour cannot have m=0
- ie no S-neighbour, since you are there
- 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
- 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 he is a survivor and has at least one S-neighbour must be proportional to $m \cdot S_{m,n}(t)$.

S: humans
n: S-neighbours
m: Z-neighbours

Moment-closure approximation

- With normalisation, this probability is then

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

- Thus, the mean number of Z-neighbours that each S-neighbour of a human has is

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

- This is called a *moment-closure approximation*
- ie we guess the higher moments of a distribution to be consistent with the already-known lower-order moments.

S: humans
Z: zombies
n: S-neighbours
m: Z-neighbours

Other moments

- Other moments may be obtained the same way
- Let $\langle i(t)\rangle_{j,k}$ be the mean number of neighbours in state i of a vertex in state j which is itself a neighbour of a vertex in state k

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

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

- We now have enough to write ODEs for the evolution of $S_{m,n}(t)$ and $Z_{m,n}(t)$.

S: humans
Z: zombies
n: S-neighbours
m: Z-neighbours

Transitions that affect a vertex X

- X, a human, is bitten by a zombie
- A human neighbour of X is bitten by a zombie
- X, a zombie, is killed by a survivor
- A zombie neighbour of X is killed.

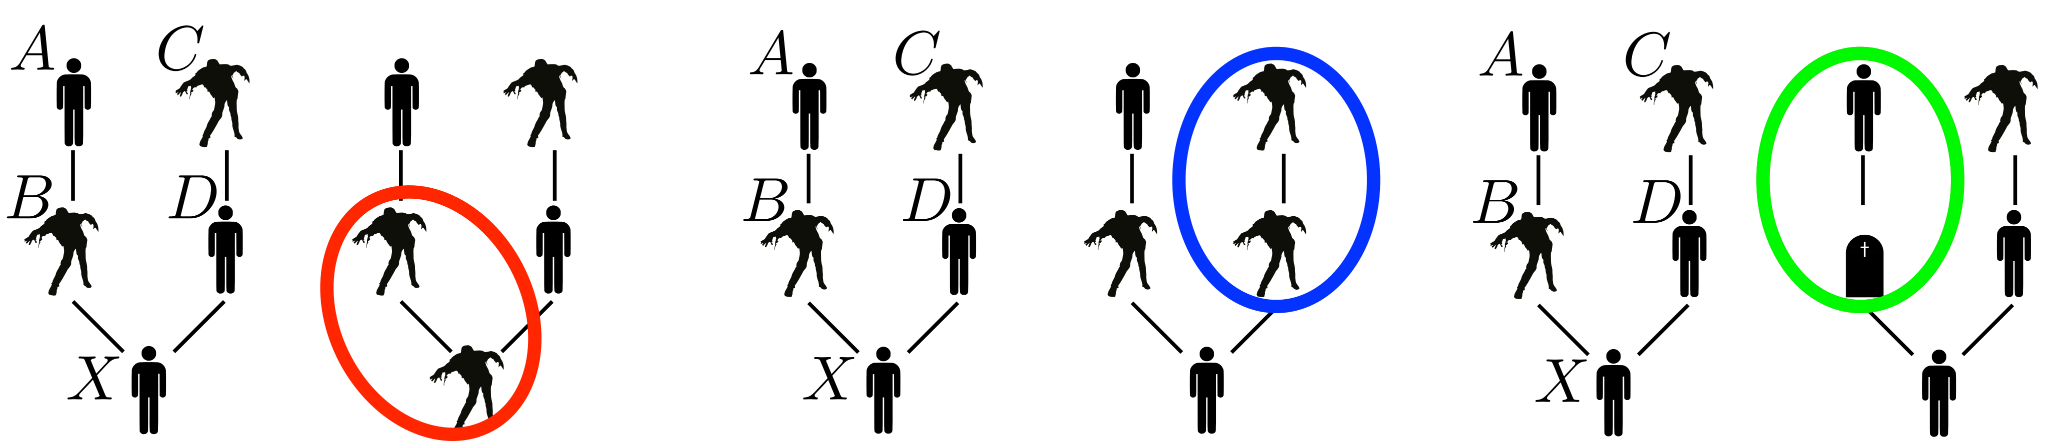

i) X, a human, is bitten by a zombie

- This causes X to become a zombie
- The transition $S_{m,n} \rightarrow Z_{m,n}$ occurs with probability $n\alpha S_{m,n}dt$
- In the figure, B bites X
- Thus, X changes from a human with one human and one zombie neighbour to a zombie with one human and one zombie neighbour.

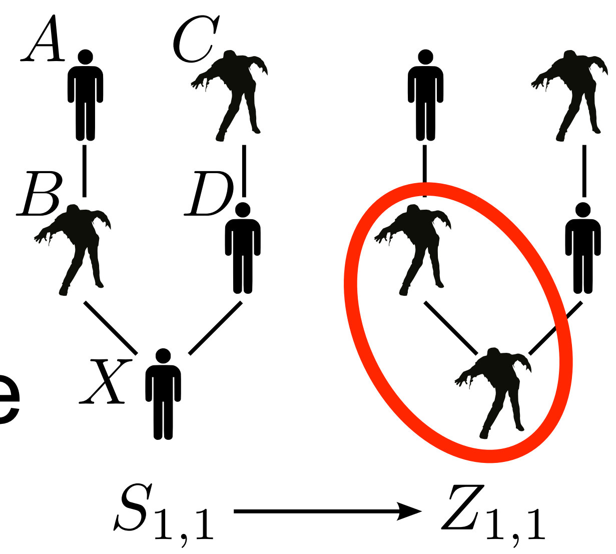

S: humans
Z: zombies
α: biting probability
n: S-neighbours
m: Z-neighbours

ii) A human neighbour of X is bitten

- This is a side effect of the previous case
- The state of X changes when one of its S-neighbours becomes a zombie
- $S_{m,n} \rightarrow S_{m-1,n+1}$ occurs with probability $m\alpha\langle z\rangle_{s,s}S_{m,n}dt$ if X is a survivor
- $Z_{m,n} \rightarrow Z_{m-1,n+1}$ occurs with probability $m\alpha\langle z\rangle_{s,z}S_{m,n}dt$ if X is a zombie
- We need to know $\langle z\rangle_{s,z}$, the average number of Z-neighbours that X has
- In the figure, C bites D.

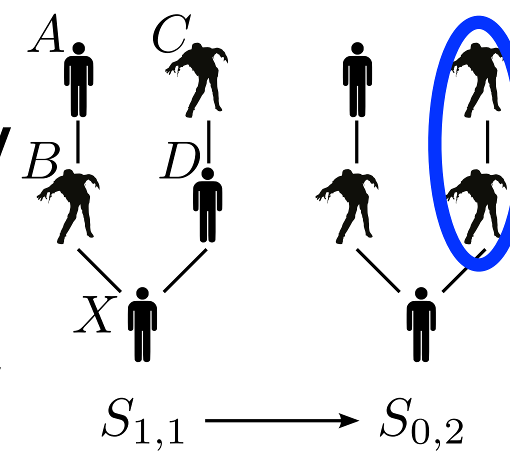

S: humans
Z: zombies
α: biting probability
n: S-neighbours
m: Z-neighbours

iii) X, a zombie, is killed by a human

- This causes X to become removed through the transition $Z_{m,n} \rightarrow R$ with probability $m\beta Z_{m,n}dt$
- Since we do not explicitly track removed individuals, we simply decrease the population of $Z_{m,n}$
- In the figure, A kills B
- This will also have consequences for X.

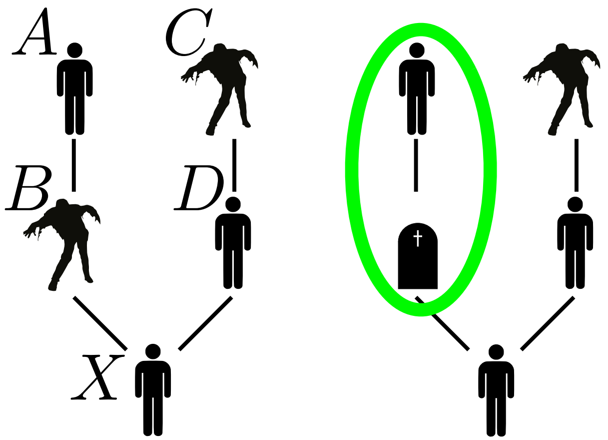

Z: zombies R: dead
β: killing probability
n: S-neighbours
m: Z-neighbours

iv) A zombie neighbour of X is killed

- This is a side-effect of the previous event
- $S_{m,n} \rightarrow S_{m,n-1}$ occurs with probability $n\beta\langle s\rangle_{z,s}S_{m,n}dt$ if X is a survivor
- $Z_{m,n} \rightarrow Z_{m,n-1}$ occurs with probability $n\beta\langle s\rangle_{z,z}Z_{m,n}dt$ if X is a zombie
- In the figure, A kills B
- This changes the number and type of neighbours that X has.

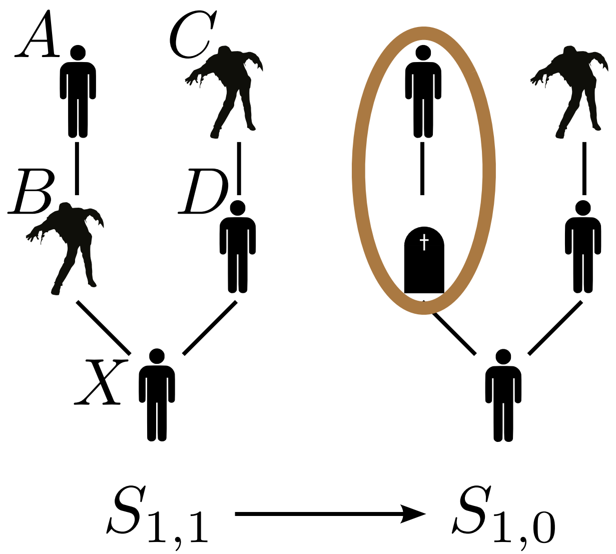

S: humans Z: zombies
β: killing probability
$\langle s\rangle_{i,j}$: mean # of neighbours
n: S-neighbours m: Z-neighbours

The ODE model

- We thus have the following ODE model

$$\frac{d}{dt}S_{m,n} = \langle z \rangle_{s,s} \cdot \alpha \left[(m+1) S_{m+1,n-1} - m S_{m,n} \right]$$
$$+ \langle s \rangle_{z,s} \cdot \beta \left[(n+1) S_{m,n+1} - n S_{m,n} \right] - n \alpha S_{m,n}$$
$$\frac{d}{dt}Z_{m,n} = \langle z \rangle_{s,z} \cdot \alpha \left[(m+1) Z_{m+1,n-1} - m Z_{m,n} \right]$$
$$+ \langle s \rangle_{z,z} \cdot \beta \left[(n+1) Z_{m,n+1} - n Z_{m,n} \right] + \alpha n S_{m,n} - m \beta Z_{m,n}$$

- Note that the ODEs are for the types of vertices as well as the number of neighbours each has.

S: humans Z: zombies
α: biting probability
β: killing probability
$<s>_{i,j}$: mean # of neighbours
n: S-neighbours m: Z-neighbours

Initial conditions

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

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

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

- 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.

S: humans
Z: zombies
n: S-neighbours
m: Z-neighbours

Choosing parameters

- Let's consider a small city of $N=10{,}000$ people
- Suppose 1% of them are already zombies, none of whom have yet been killed
- ie $\epsilon=0.01$
- Suppose everybody has equal chance to be connected to everybody else
- Furthermore, we know λ, the mean number of acquaintances that people have
- We can then determine $\{p_k\}$ by a Poisson distribution for large populations.

ϵ: initial zombie fraction
p_k: degree distribution

Poisson distribution

- The Poisson distribution applies to large populations when everybody has similar numbers of acquaintances
- In this case, the degree distribution is

$$p_k = \frac{\lambda^k e^{-\lambda}}{k!}$$

- Thus, we can determine everything from the mean degree
- But what happens if people have vastly different number of acquaintances?
- A few very popular people may make an enormous difference.

λ: mean # acquaintances
k: degree
p_k: degree distribution

The power law

- Some people are just more popular than others
- Or, more precisely, they come into contact with more people than the rest of us
- The *power-law* distribution describes this

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

- τ is a scalable parameter
- We truncate at k_{max}.

k: degree
p_k: degree distribution

Two degree distributions

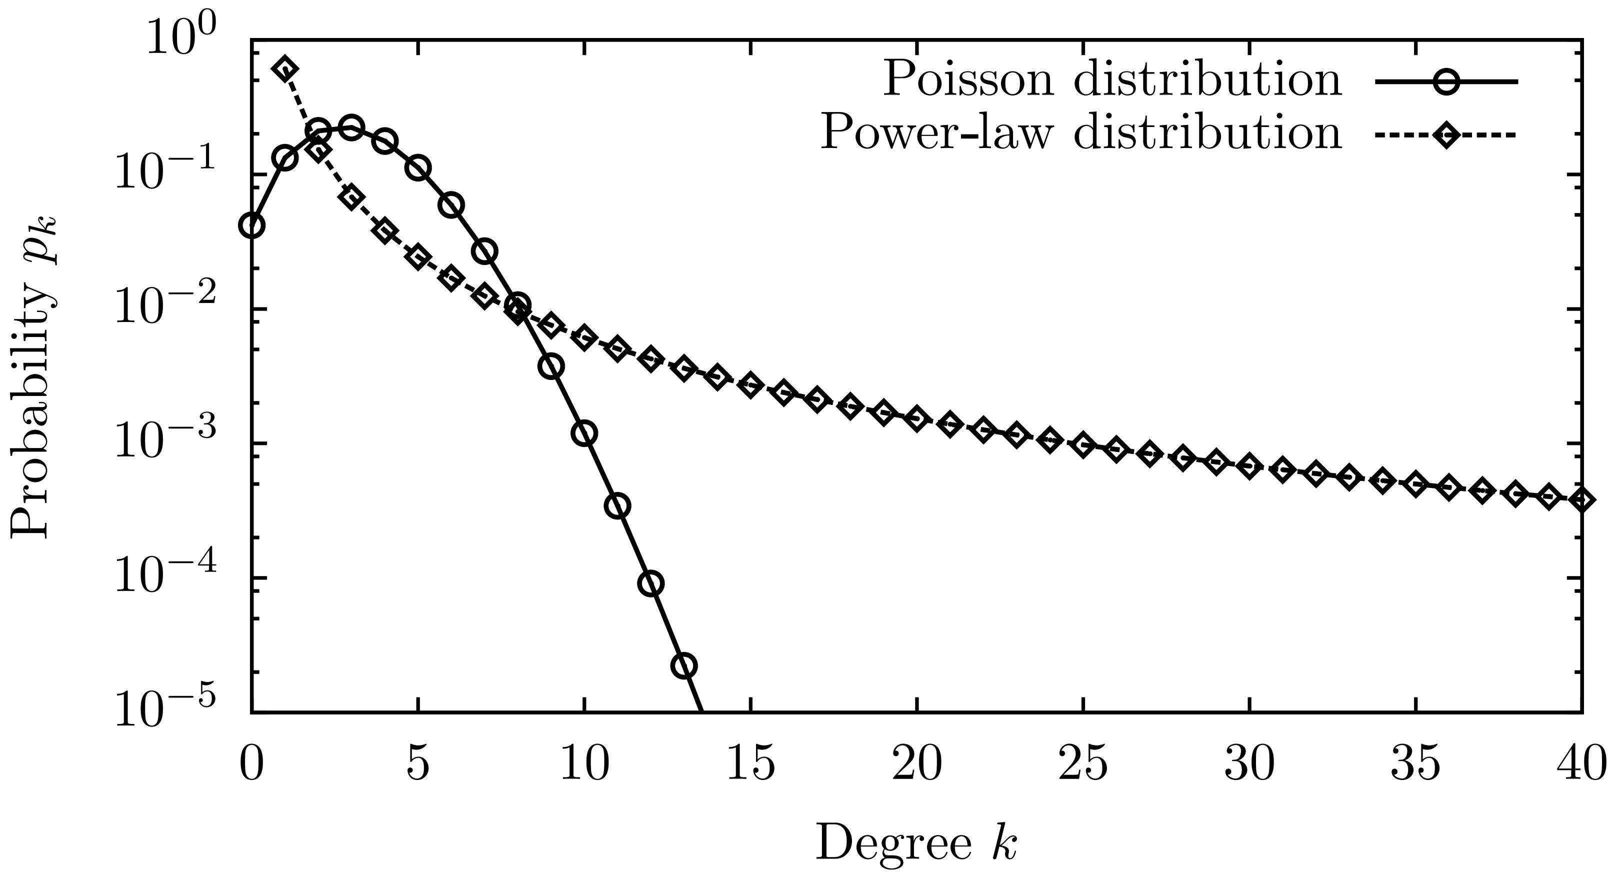

- Same average degree λ=3.17
- For the power law, τ=2, k_{max}=100
- Some people have a *lot* of friends.

λ: mean # acquaintances
k_{max}: max degree
τ: scalable parameter

Our two approaches

- We have two approaches
  - Monte Carlo simulations and integrating ODEs
- Let's apply these to two populations that differ only by their degree distribution
- Curves show results for ODEs
- Symbols show results of averaging over 100 Monte Carlo simulations
- We also illustrate the final state for each
- These are done using GUESS software with an algorithm that distinguishes vertices by their state (colour) and degree (size).

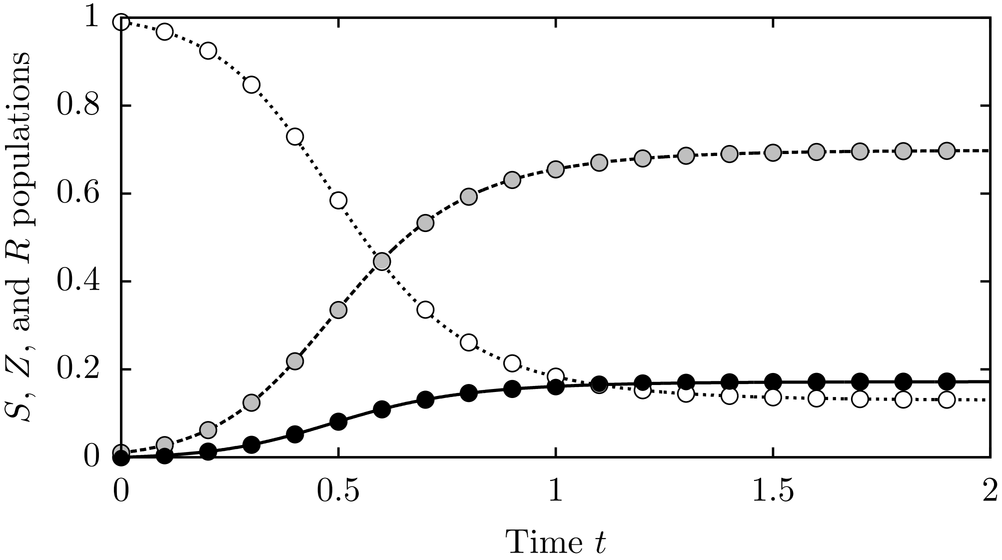

(a) Evolution for Poisson distribution.

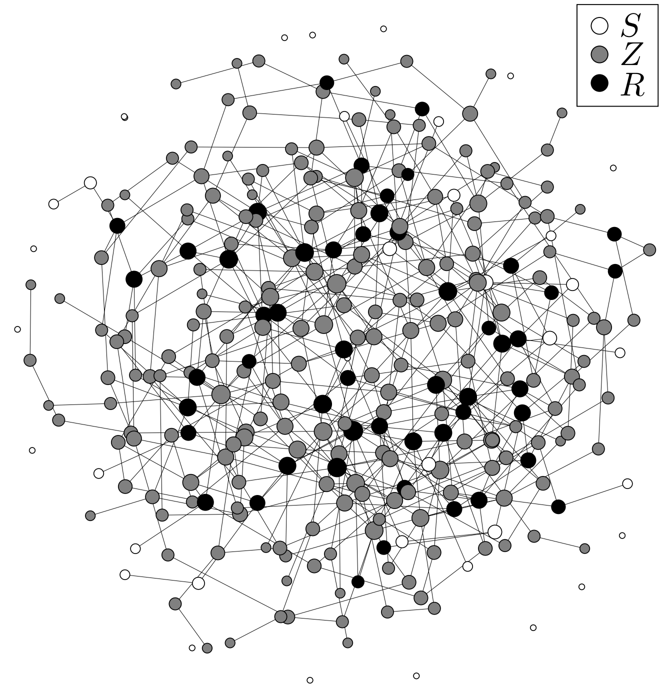

(b) Final state for Poisson.

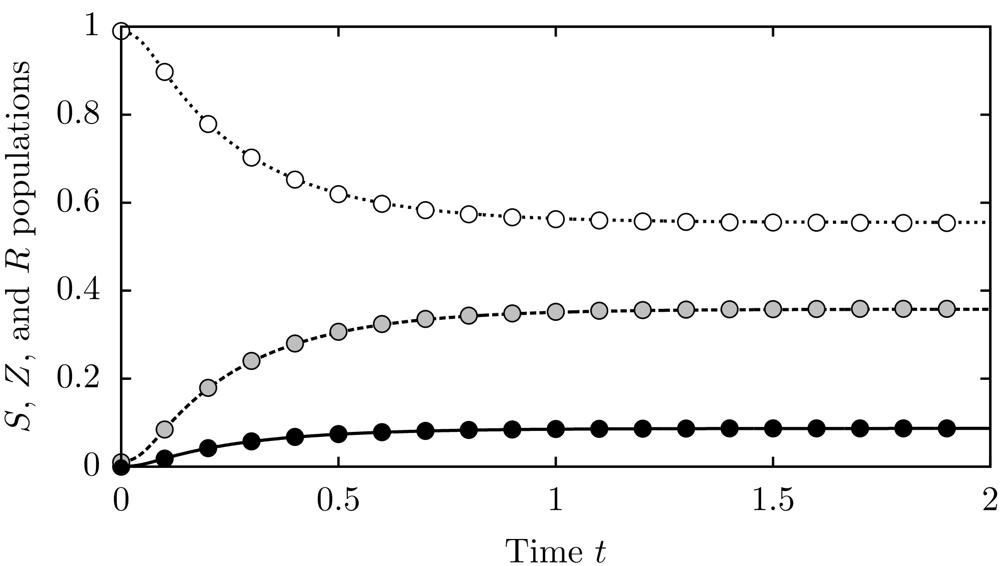

(c) Evolution for power-law distribution.

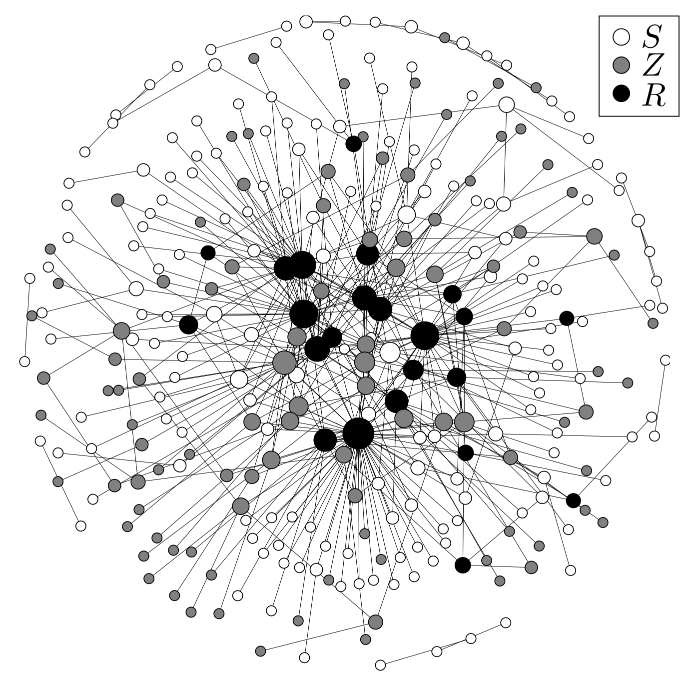

(d) Final state for power-law.

Similarities between the methods

- The Monte Carlo simulations and ODE solutions have striking similarities
- When a sufficiently small time increment Δt is used for the Monte Carlo simulations, the results should be arbitrarily close to an exact realisation of the rules we chose to model
- Thus, if the integrated system of ODEs agrees with the simulations, then they in turn accurately reflect the chosen rules.

Testing the assumptions

- This accuracy means that the information ignored in building the ODEs does not have much impact on the propagation dynamics of the whole system
- Thus, everything of importance has been included in the model
- This is a very powerful result.

Degree distribution differences

- There are also significant differences between the results for the two degree distributions
- In a non-network model, these choices would both give the same results
- Because they lead to identical average degree for vertices
- This justifies the need for a network-based approach to the problem.

Coexistence with the undead

- No equilibrium can exist in the model as long as there are humans with zombie neighbours
- Either the human would eventually kill the zombie or the zombie would bite the human
- Hence, humans and zombies can coexist...

  ...but only if zombies are clumped and humans are clumped
- That is, the only paths between humans and zombies go through removed vertices
- The dead form a barrier.

Managing the zombie invasion

- We should develop very efficient tools to eliminate zombies at the outset
- These dead zombies will also protect us from undead zombies
- Notice too that things are actually better with the power-law degree distribution
- The presence of high-degree vertices, aka *hubs*, usually makes diseases worse
  - eg sex workers for STIs
  - hospitals or schools for pulmonary infections
- Hubs facilitate fast disease transmission.

Zombie hubs

- The difference with zombie hubs is that humans can also obliterate them
- That is, if humans take out an entire zombie hub, they seriously slow down the transmission
- These removals reduce the average degree of the remaining vertices, since their degree was originally large compared to the mean
- In a more “balanced” situation, like a Poisson distribution, hubs are mostly absent
- Thus, zombie death does not have this knock-on effect.

Adapting for realism

- There are other factors we can add in
- Rather than fight to the death, some humans will flee when the situation becomes too dire
- Also, the undead are usually seen in hordes
- We have only considered *static* networks
  - ie the same edges always linked the same vertices
- We can now add in *adaptive networks*
  - the structure of the network itself is now allowed to change
  - this comes at an additional cost, however.

The new rules of the game

- During an infinitesimal time interval [t,t+dt), a human has probability γdt to flee from a zombie
- This human will stop at the next individual she encounters
- The human disconnects from the zombie and reconnects to either a human ($S_{m,n} \rightarrow S_{m+1,n-1}$) or a zombie ($S_{m,n} \rightarrow S_{m,n}$)

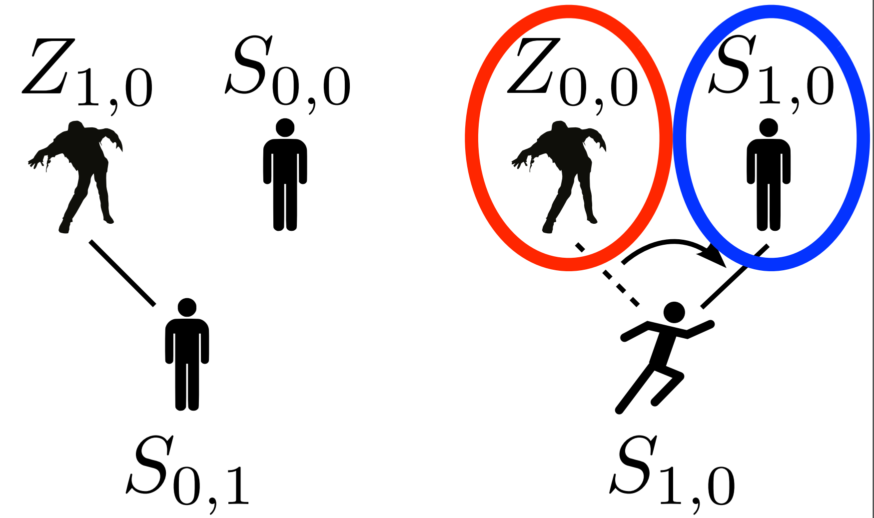

- This causes "side effect transitions" to both the old and new neighbour.

S: humans
Z: zombies
n: S-neighbours
m: Z-neighbours

Adapting the rules

- An individual is better at attacking the enemy when she is helped by neighbours of her own type (S or Z)
- During the interval $[t,t+dt)$, a zombie with n fellow flesh-eaters can bite a human with probability $\alpha(1+hn)dt$
- A human defending herself alongside m friends will kill a zombie with probability $\beta(1+hm)dt$
- The H-factor h determines how useful friends are in combat.

S: humans Z: zombies
α: biting probability
β: killing probability
n: S-neighbours m: Z-neighbours

Average number of Z-neighbours

- For notational convenience, define the average number of Z-neighbours a human has:

$$\langle z \rangle_s = \sum_{m,n} n S_{m,n}$$

- The probability for any individual to encounter a human running away from a zombie depends on this quantity.

S: humans Z: zombies
$<s>_{i,j}$: mean # of neighbours
n: S-neighbours m: Z-neighbours

The adaptive system

- We thus have

$$
\begin{aligned}
\frac{d}{dt}S_{m,n} = {} & \left(h\langle z\rangle_{z,s}+1\right)\langle z\rangle_{s,s}\alpha\left[(m+1)S_{m+1,n-1}-mS_{m,n}\right] \\
& + \left(h\langle s\rangle_{s,z}+1\right)\left(\langle s\rangle_{s,z}-1\right)\beta\left[(n+1)S_{m,n+1}-nS_{m,n}\right] \\
& + (1+hm)\,\beta\left[(n+1)S_{m,n+1}-nS_{m,n}\right]-\left(h\langle z\rangle_{z,s}+1\right)n\alpha S_{m,n} \\
& + \gamma\frac{S}{S+Z}\left[(n+1)S_{m-1,n+1}-nS_{m,n}\right]+\gamma\langle z\rangle_{s}\frac{S_{m-1,n}-S_{m,n}}{S+Z} \\
\frac{d}{dt}Z_{m,n} = {} & \left(h\langle z\rangle_{z,s}+1\right)\left(\langle z\rangle_{s,z}-1\right)\alpha\left[(m+1)Z_{m+1,n-1}-mZ_{m,n}\right] \\
& + \left[(n-1)h+1\right](m+1)\alpha Z_{m+1,n-1}-(hn+1)\,m\alpha Z_{m,n} \\
& + \left(h\langle s\rangle_{s,z}+1\right)\langle s\rangle_{z,z}\beta\left[(n+1)Z_{m,n+1}-nZ_{m,n}\right] \\
& - \left(h\langle s\rangle_{s,z}+1\right)m\beta Z_{m,n}+\left(h\langle z\rangle_{z,s}+1\right)n\alpha S_{m,n} \\
& + \gamma\left[(m+1)Z_{m+1,n}-mZ_{m,n}\right]+\gamma\langle z\rangle_{s}\frac{Z_{m-1,n}-Z_{m,n}}{S+Z}
\end{aligned}
$$

Varying the H-factor

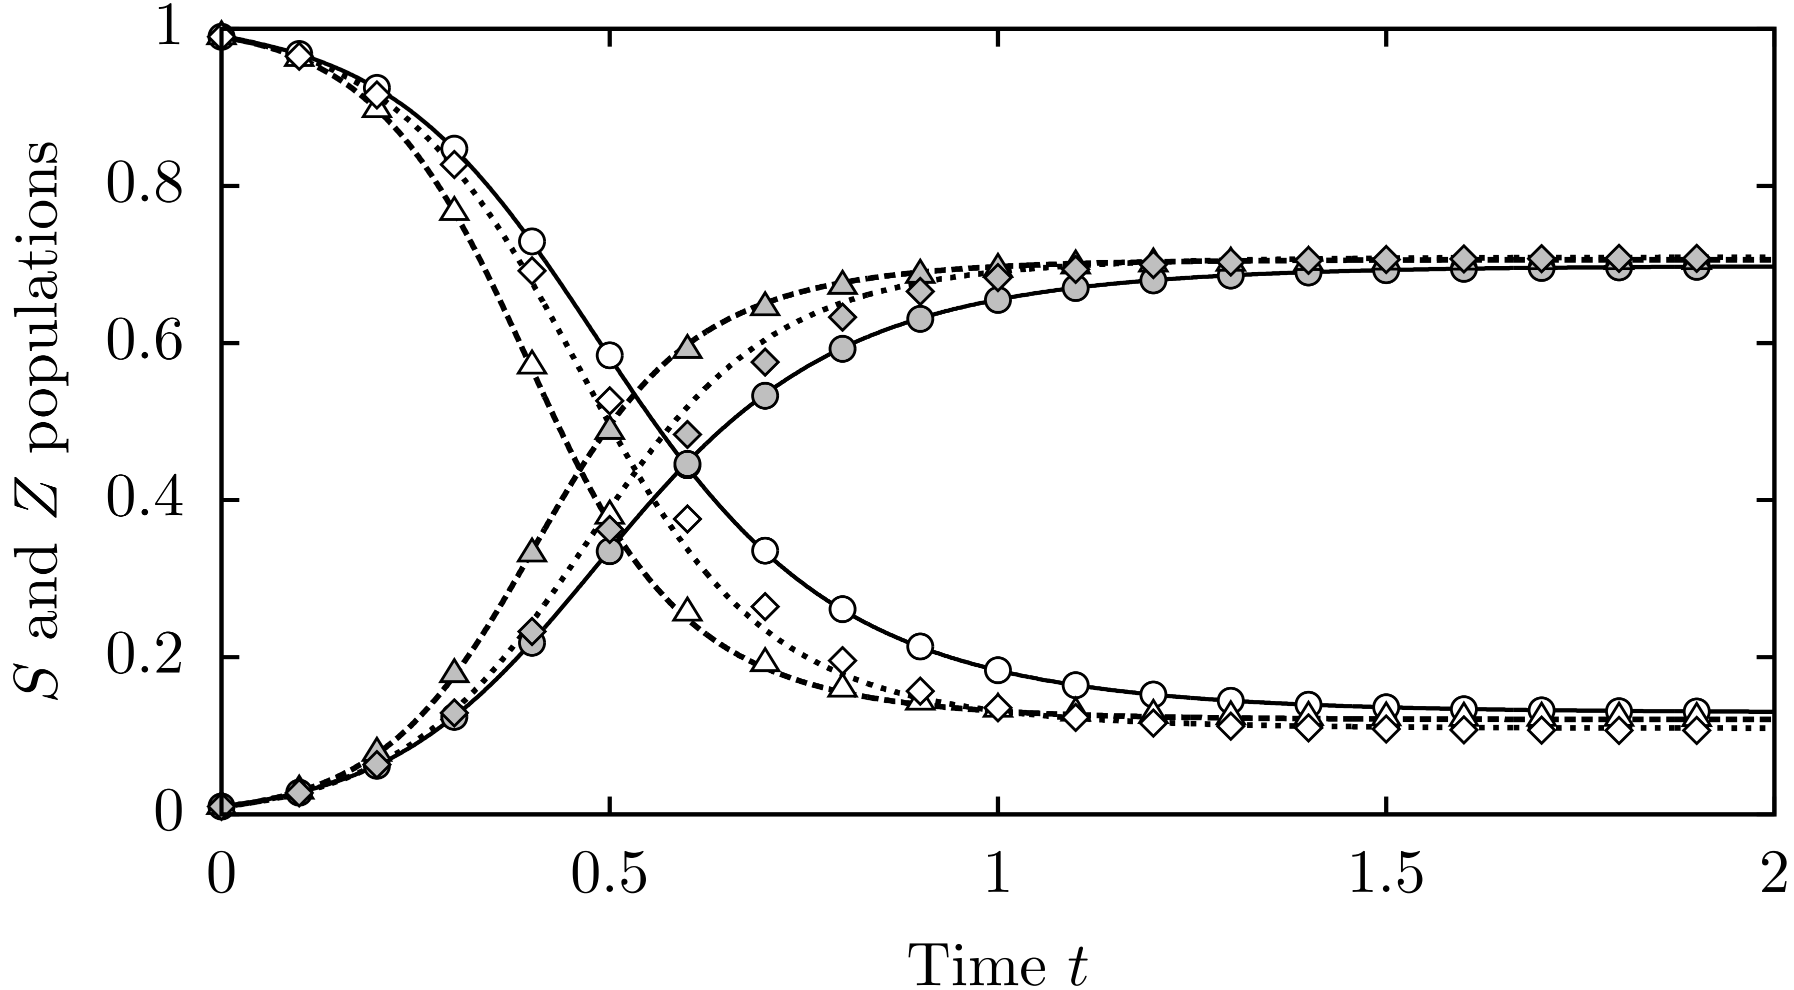

○ h=γ=0 (nobody helps, nobody flees)

△ h=λ^{-1}, γ=0 (help, nobody flees)

◇ h=λ^{-1}, γ=2 (help, but people flee).

h: helping factor
λ: mean # acquaintances

Zombies like it when your friends help

- We expect it to be the other way around...

  ...but only if humans are more efficient at killing zombies than they are at biting
- Allowing humans to flee from zombies gives them better life expectancy...

  ...but it only postpones the inevitable
- All scenarios result in only 15% human survival
- Thus, humans should eradicate zombies early, because we can’t survive a zombie onslaught, even with help from our friends.

Including flight matters

- The fit between the Monte Carlo simulations and the analytical predictions is less satisfying when flight is possible
- Thus, something we either approximated or ignored when obtaining the ODEs becomes important when flight is allowed
- The most likely culprit is correlation.

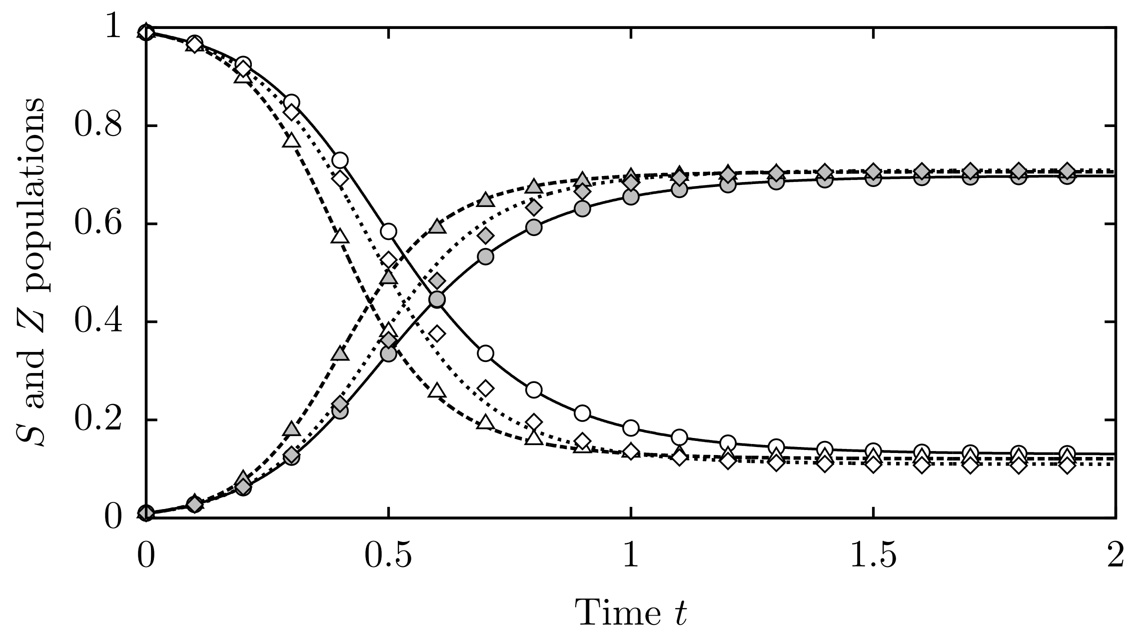

The average survivor

- In our ODEs, when a human flees from a zombie and then encounters vertex X, we add a new S-neighbour to vertex X
- Later on, we consider this human to be the *average survivor*
- This is not strictly true
- She is not just the average survivor, but the average survivor who just ran away from a zombie
- This additional information tells us something about her state.

A built-in correlation

- Knowing the average survivor has fled from a zombie tells us something
- She would unlikely to have made it this far had she been in contact with 100 zombies
- Nor would she be fleeing had she been in contact with at least one zombie originally
- The fact that our ODEs do not consider these correlations is probably the principal cause of the disagreement with the Monte Carlo simulations.

Including neighbours of neighbours

- We could adapt for this information
- However, this would require including the state of the neighbours of the neighbours of each vertex

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

where each $m_{i,j}$ represents the number of human neighbours who themselves have i human and j zombie neighbours

- And similarly for $n_{i,j}$.

S: humans Z: zombies
n: S-neighbours m: Z-neighbours

Exploding complexity

- Since a differential equation would be required for each of these quantities, it's easy to see how quickly the complexity of the model would explode
- However, considering the error observed in the results was not that critical in the first place, the simpler approach is sufficient
- Unlike zombies, we can live with this.

Network summary

- Network modelling allows us to include a topology in a model of interaction between a large number of elements
- They try to account for mechanistic behaviours at the individual level
- The interactions then extend to the bigger picture
- Thus, networks are useful when considering the effects of local dynamics on the global evolution of the system.

Effects we considered

- Heterogeneity of behaviours
  - (some individuals can be disconnected from the rest of the world, while others might end up in the middle of the zombie apocalypse)
- Local environment
  - (individuals are only affected by their immediate environment, allowing a single zombie to start a virulent, but localised outbreak)
- Social behaviours
  - (humans can flee from zombies and regroup, while the undead can hunt in hordes).

What else could we have included?

- Types of individuals
  - not all humans will react the same in a zombie invasion
  - eg the behaviour of a child differs from that of a soldier
  - network modelling is well-suited to considering such heterogeneity
- Social structure
  - we gave very basic behaviour to both the living and the dead
  - however, the formation of more complex social structure could be taken into account.

The cost of including more realism

- Network modelling is a natural approach to the description of individual and heterogeneous behaviour
- However, the inclusion of such structure usually complicates its treatment
- Furthermore, complexity increases rapidly in network theory.

Zombie summary

- Any invasion must be stopped very early
- Otherwise, a significant fraction of the population will be either dead or zombified
- This highlights the importance of developing efficient and powerful anti-zombie defences
- Including cooperation actually accelerated the zombie invasion (if $\alpha>\beta$)
- Flight only slows down the progression of the invasion, without modifying its outcome
- Our best hope is training so that $\alpha<\beta$
- ie efficient zombie fighting skills

α: biting probability
β: killing probability

Roaming, isolated zombies

- Within the context of our model, an invasion can be stopped, but not reversed
- Once a barricade is established, we can expect to find groups of roaming, isolated zombies
- Wiping out these groups may require further interventions
- eg weapons of mass destruction
- This should only be done once the system has stabilised.

Authors

- Laurent Hérbert-Dufresne (Université Laval)
- Vincent Marceau (Université Laval)
- Pierre-André Noël (Université Laval)
- Antoine Allard (Université Laval)
- Louis J. Dubé (Université Laval)

L. Hérbert-Dufresne, V. Marceau, P.-A. Noël, A. Allard, L.J. Dubé, The Social Zombie: Modelling undead outbreaks on social networks. (In: R. Smith? (ed) Mathematical Modelling of Zombies, University of Ottawa Press, *in press*.)