Applied Computational Mathematics Lab

EPIDEMIC MODELING · MODULE 4

Spatio-temporal models

Patches linked by travel, networks of people, waves in continuous space

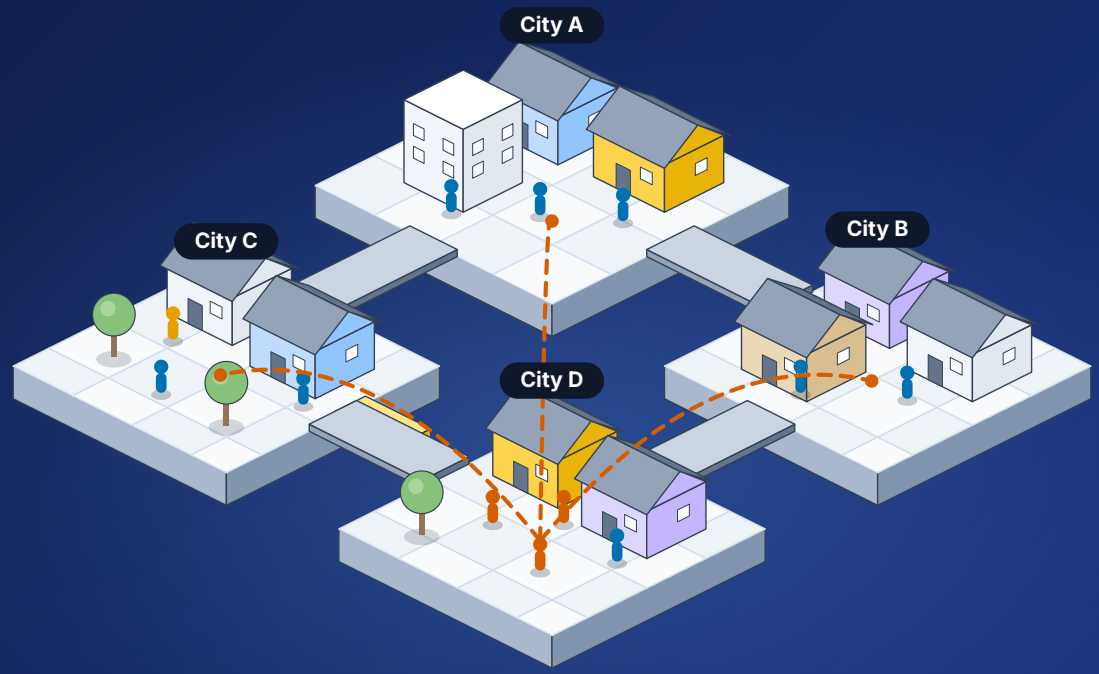

4 EPISODES

1. Cities linked by travel: metapopulation models
2. Networks of people
3. Epidemics in continuous space
4. Bounding the spread: norms and stability

Based on Chapter 4 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)

acml.teddylazebnik.com

BEFORE WE START

Warm-up: true or false?

CHECK YOURSELF

(a) In a chain of cities, more travel means more people infected in the end.

(b) Two contact networks with the same average number of contacts have the same epidemic threshold.

(c) In a reaction–diffusion model, four times more movement makes the wave twice as fast.

(d) If the spectral radius of the spread matrix is below one, the total number of cases can never grow.

HOW TO USE THIS

Guess now; each part of the deck settles one statement, and the answers are at the end.

WHY THIS MODULE

Module 3 assumed everyone mixes with everyone. This module puts people in space: in places, on contact networks and on maps.

BEFORE WE START

This module's map

PART 1

Patches linked by travel

Metapopulations, two kinds of coupling, timing, effective distance, imported sparks, R_0 of a network.

PART 2

Networks of people

Transmissibility, percolation, excess degree, hubs, clustering, interventions on the graph.

PART 3

Continuous space

Reaction–diffusion, wave speed, borders, drift, uneven landscapes, long jumps.

PART 4

Norms and stability

Bounds on growth, transient waves, simple policy rules, from space back to graphs.

GOALS

By the end you should be able to build a metapopulation model, compute an epidemic threshold on a network, estimate a wave speed, and use matrix norms to bound growth.

ONE WORD

Spatio-temporal model: the state depends on where as well as when: per patch, per person on a network, or per point of a map.

Applied Computational Mathematics Lab

PART 1

Patches linked by travel

Each city runs its own SIR; travel couples them.



Flux form: people move

$$\frac{dS_p}{dt} = -\beta_p \frac{S_p I_p}{N_p} + \sum_{q \neq p} (m_{qp} S_q - m_{pq} S_p)$$

$$\frac{dI_p}{dt} = \beta_p \frac{S_p I_p}{N_p} - \gamma_p I_p + \sum_{q \neq p} (m_{qp} I_q - m_{pq} I_p)$$

- Patches p = 1, ..., P (cities, districts); each runs SIR with its own rates.
- m_{pq}: per-capita travel rate from p to q. Every compartment moves, so an infected traveller arrives infected.
- Movement only relocates people: the total over all patches is conserved.

THE MATRIX M

A travel matrix with zero diagonal. Read across a row: where people of patch p go. Read down a column: where the people arriving in q come from.

WHEN TO USE IT

Long stays, migration, people displaced by a crisis: the flux form is natural when people really relocate.

TRY IT

Three patches, all rates 0.01 per day. How many people leave patch 1 per day if it has 50,000 people? Where do they go?

Mixing form: risk moves

$$\lambda_p(t) = \sum_q \beta_{pq} \frac{I_q(t)}{N_q(t)}, \qquad \frac{dS_p}{dt} = -\lambda_p S_p$$

- People keep a home patch. Infection pressure travels: commuters and visitors meet infected people elsewhere and come home.
- β_{pq} bundles how much time people of p spend with people of q, and how infectious that contact is.
- Closer to what phone and travel data measure: time spent elsewhere, origin–destination flows, shared venues.

FLUX

people relocate

state counts are who is present now; fits migration and long stays.

MIXING

risk is shared

state counts are residents; fits commuting and visits.

IN PLAIN WORDS

A common slip: reading the flux model's I_p as "residents of p". In the flux form it is "people in p right now".

Two clocks: growth and travel

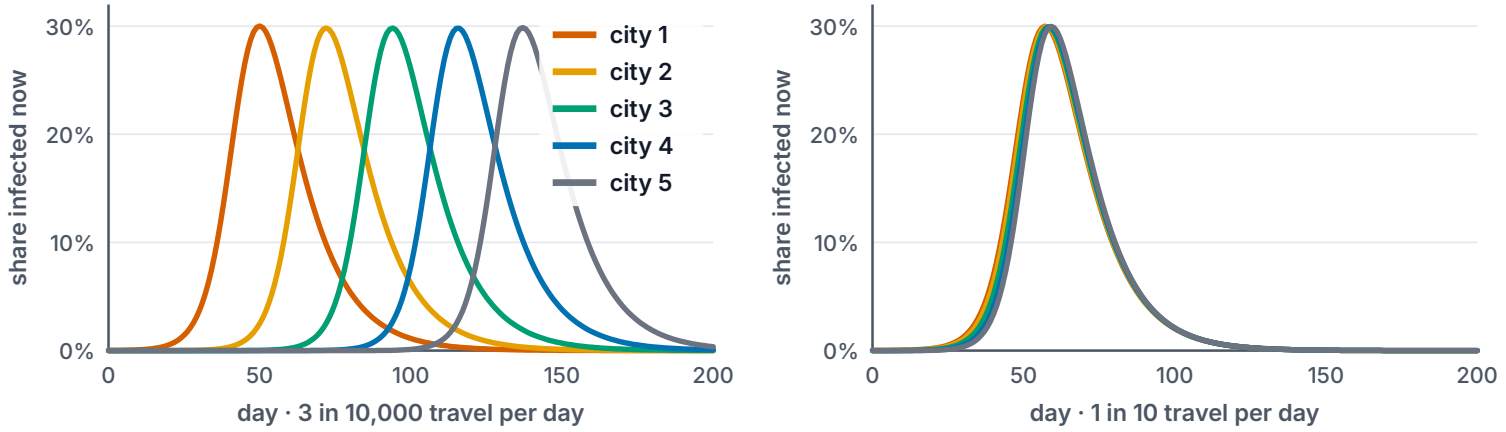

Our runs: five cities in a row (R_0 = 3 each), neighbours linked by travel, 1 in 10,000 infected in city 1 on day 0.

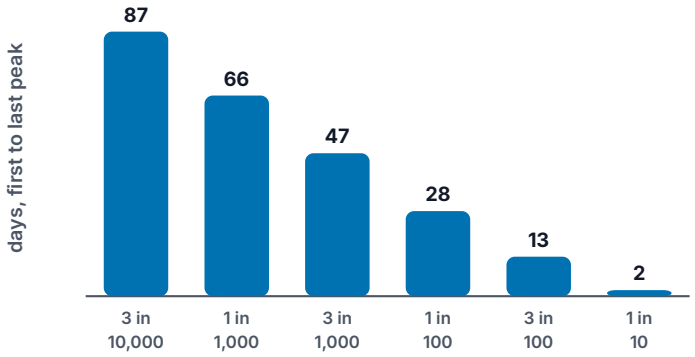

- Weak travel: a chain of local waves; the last city peaks 87 days after the first.
- Strong travel: the cities synchronize and behave like one well-mixed town.
- Either way each city ends with about 94% infected: travel moves the timing, not the total.

Effective distance

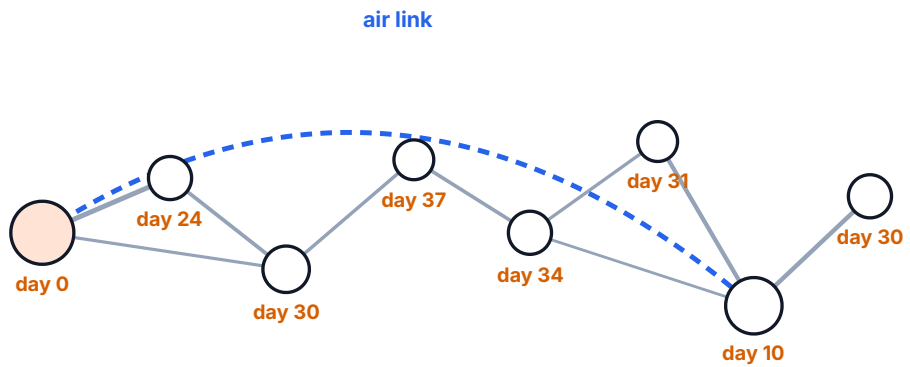

Our toy map: eight cities, travel falling with distance cubed (scaled by destination size), plus one air link. Labels: arrival day (1 in 10,000 infected).

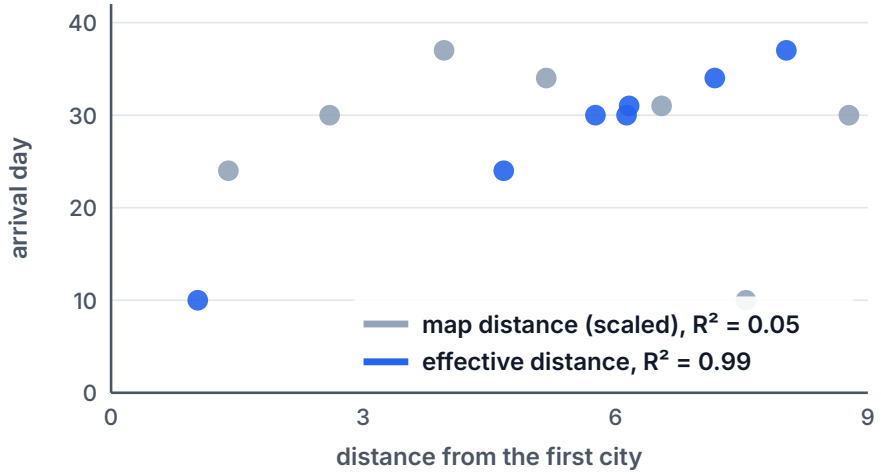

- The air-linked city is reached on day 10, before the nearest neighbour (day 24).
- Map distance explains almost nothing: $R^2 = 0.05$.
- Effective distance (Brockmann and Helbing, 2013): $d_{pq} = 1 - \ln P_{pq}$, where P_{pq} is the share of travellers leaving p who go to q, added up along the shortest route.
- Against effective distance, arrival times fall on a line: $R^2 = 0.99$.

IN PLAIN WORDS

Contagion spreads like a wave in travel space, not on the map: busy long-range links decide arrival times.

Imported sparks

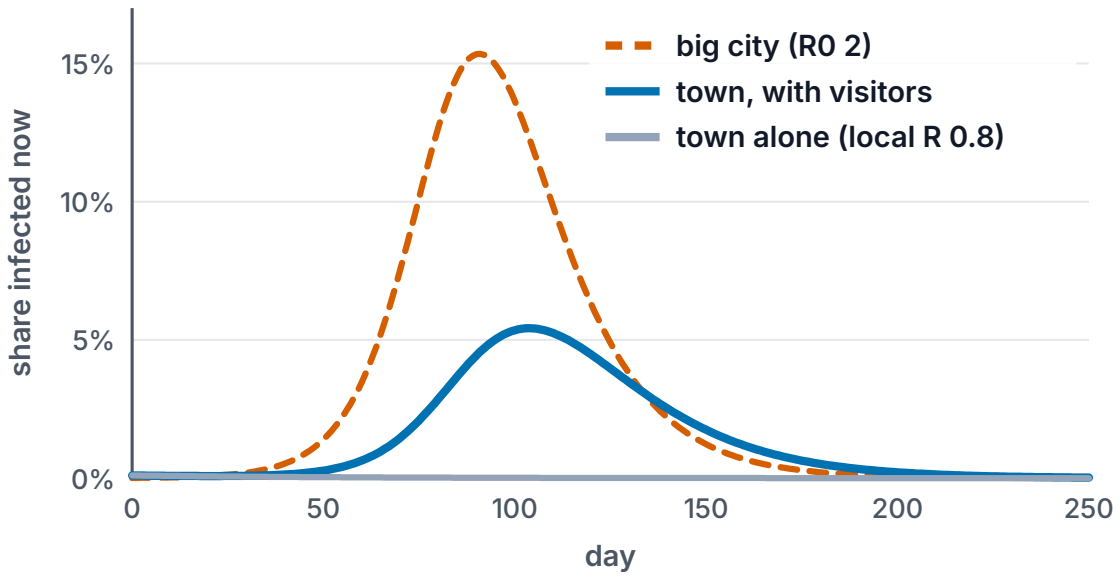

Our run: the town's force of infection adds 10% of the big city's infection pressure (visitors).

- Alone, the town (local R = 0.8) lets a few cases fizzle: 0.5% are ever infected.
- Linked to a big city with $R_0 = 2$, it ends with 36% infected, though locally R stays below one.
- With chance in the model, small places fade out and get reseeded again and again: persistence comes from the balance of local extinction and reintroduction.

IN PLAIN WORDS

Ask two questions: is R above one locally, and how often do sparks arrive?

R_0 of a network of patches

$$R_0 = \rho(\mathrm{diag}(R_1, \ldots, R_P)\, M), \qquad M = (1-c)\, I + \frac{c}{P}\, J$$

- The next-generation matrix of Module 3, with patches as types: local reproduction numbers R_p, combined by a mixing matrix (row p: where people of p spend their contacts).
- Our example: one hot city (local R = 1.6) and four cool ones (0.8), equal sizes; a share c of contacts is spread evenly over all five.
- Mixing dilutes the hotspot: R_0 falls below one once c reaches 0.75; full mixing gives the average, 0.96.

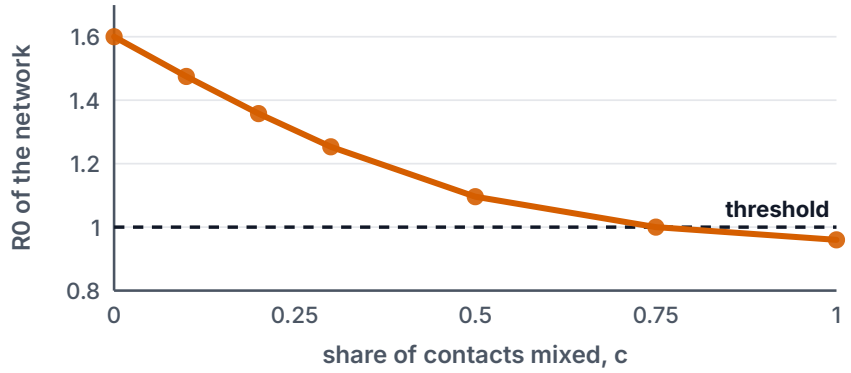

MIXING C	0	0.2	0.5	0.75	1
R_0	1.60	1.36	1.10	1.00	0.96

THE OTHER DIRECTION

Flip it: five cities at 0.95 (R_0 = 0.95). Make one a hotspot at 1.6, with c = 0.3: R_0 = 1.27. One hotspot can carry a region.

TRY IT

With full mixing and equal patch sizes, why is R_0 the average of the local numbers?

Extensions, and a border closure

- **Different patches**: own β, γ and interventions; some patches export infection (sources), others import it (sinks).
- **Commuting**: day and night, weekday and weekend mixing differ; a time-varying $M(t)$ or sub-compartments for people away from home.
- **Changing travel**: holidays, crises and restrictions; letting M change avoids blaming behaviour on biology.

CLOSURE (TRAVEL ÷ 20)	ARRIVAL IN B	PEAK IN B	FINAL SIZE IN B
none	day 16	day 61	94%
on day 20	day 16	day 64	94%
on day 7	day 20	day 68	94%
on day 0	day 28	day 75	94%

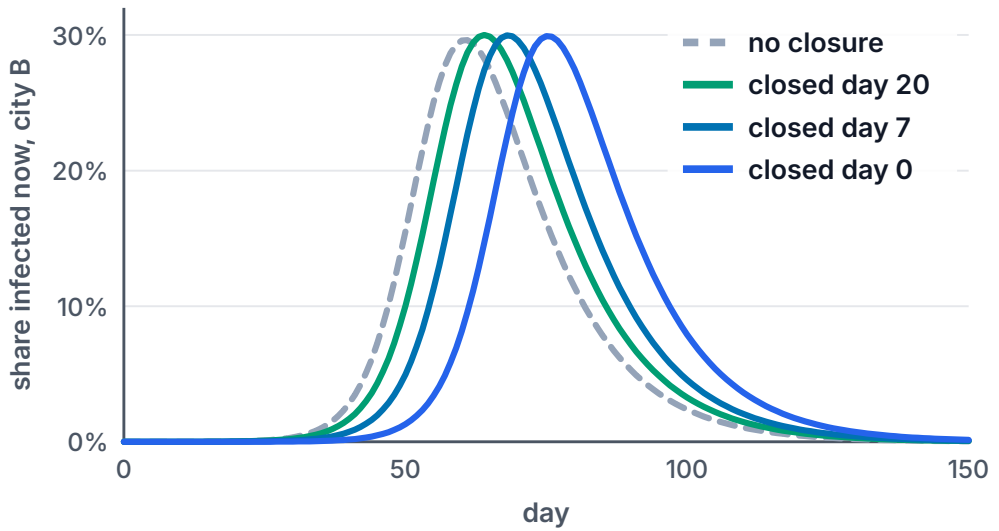

Our runs: two cities, travel 0.003 per day, cut to 1/20 on the day shown.

IN PLAIN WORDS

A closure buys time to prepare, treat and vaccinate; without those, the final size is unchanged. Closing after arrival buys little.

Check yourself: Part 1

CHECK YOURSELF

1. Two cities, each with $R_0 = 2.5$, are linked by a little travel. If travel grows tenfold, what changes in the second city: its final size, or when its wave comes?

2. A town's local R is 0.9. Give two ways it can still see many cases.

3. Five equal patches with local R = 2, 0.5, 0.5, 0.5, 0.5 mix fully. What is R_0, and does the outbreak grow?

TIP

Answers are at the end of the deck. Try them before you look.

Applied Computational Mathematics Lab

PART 2

Networks of people

Who is connected to whom.



Contact networks and T

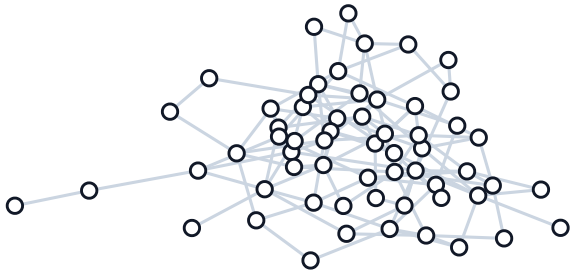

A drawing of 60 people, about 4 contacts each (illustration).

$$T = \frac{\tau}{\tau + \gamma}$$

- Nodes are people; an edge is a recurring chance to infect (household, desk, meeting).
- τ: transmission rate along one edge; γ: recovery rate.
- T is the chance an edge transmits before the infected end recovers: two competing exponential clocks, and τ wins with probability τ/(τ + γ).
- Our runs: τ = 0.05 and γ = 0.1 per day, so T = 1/3.

MODEL

The exact process is a continuous-time Markov chain on the states of all nodes. Simulate it event by event; it becomes the well-mixed SIR only when the network washes out.

Percolation and the threshold

$$T \cdot \kappa > 1, \qquad \kappa = \frac{\langle k^2 \rangle - \langle k \rangle}{\langle k \rangle}$$

- Keep each edge with probability T: a large outbreak needs a giant connected cluster of kept edges (bond percolation). For final sizes on random networks this picture is close to exact.
- κ is the mean excess degree: the other contacts of someone you reach by following an edge. Infection follows edges, so it meets well-connected people more often.
- Friendship paradox, in our hub-heavy network: an average person has 4.1 contacts, the person at the end of a random edge 11.5.

NETWORK (2,000 PEOPLE)	MEAN K	κ	$T_C = 1/\kappa$	CLUSTERING
wired at random	4.0	4.0	0.25	0.00
everyone 4, random	4.0	3.0	0.33	0.00
everyone 4, clustered	4.0	3.0	0.33	0.42
hub-heavy	4.1	10.5	0.10	0.01

TRY IT

Half of a town has 2 contacts and half has 6. Compute the mean degree, κ and the threshold T_C. Compare with everyone having 4.

Same contacts, different epidemics

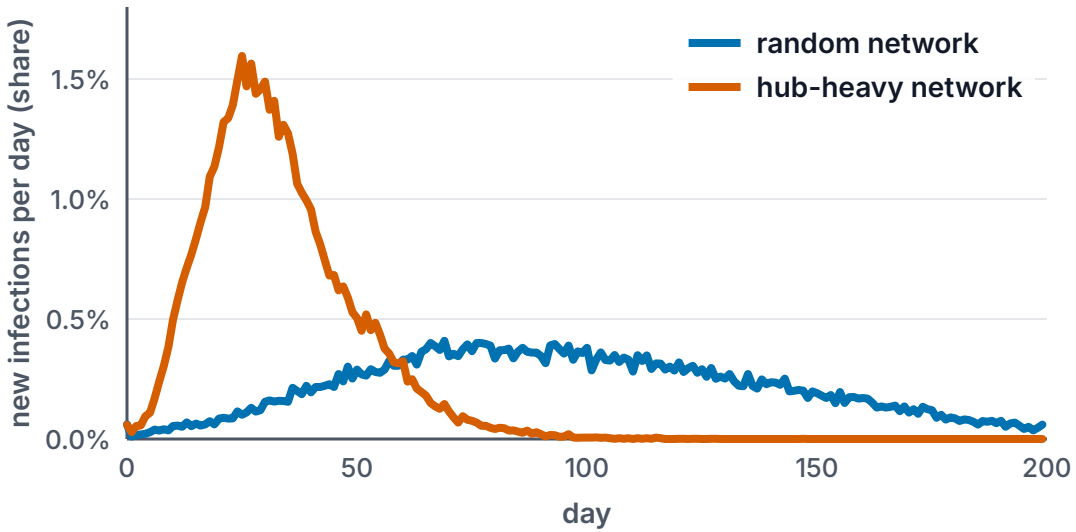

Our runs: exact event-driven SIR, T = 1/3, mean of the large outbreaks (at least 5% infected), 200 runs per network.

NETWORK	T = 1/3: TAKE OFF	SIZE	PEAK DAY	T = 1/2: TAKE OFF	SIZE
random	35%	44%	105	66%	80%
hub-heavy	36%	50%	27	63%	73%
everyone 4	11%	14%	115	69%	85%
clustered	0%	0%	–	13%	12%

- Hubs lower the threshold and speed things up: a peak on day 27 instead of 105.
- At T = 1/3 the regular network sits right at its threshold, so few runs take off.

Hubs first, and clustering

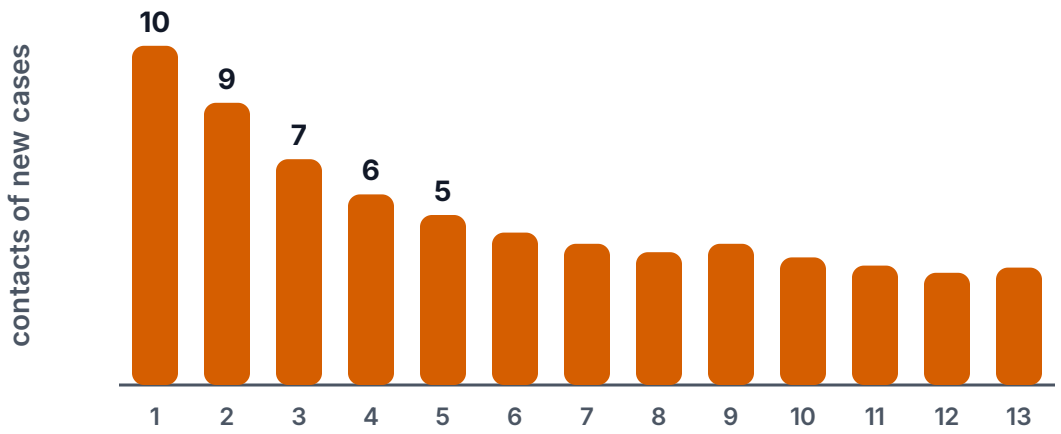

Our runs on the hub-heavy network: average contacts of people infected in each week (everyone: 4.1).

- Hubs are infected early and pass it on fast; then the epidemic runs out of hubs and slows. Early growth can overstate what follows.
- Triangles waste transmissions: by the time you could infect a friend's friend, the friend may already have done it.
- Everyone with exactly 4 contacts, T = 1/2: wired at random, 69% of runs take off and infect 85%; with neighbours sharing friends (clustering 0.42), 13% take off and infect 12%.

IN PLAIN WORDS

Our clustered network is a ring of neighbours with a few shortcuts: some of the slowdown comes from long paths, not only from triangles.

Diseases you can catch again

$$\frac{\tau}{\gamma} > \frac{1}{\rho(A)}$$

- For SIS on a fixed network, a common approximation puts the persistence threshold at one over the largest eigenvalue of the adjacency matrix A.
- Dense cores and tightly knit groups raise ρ(A): about 13 for our hub-heavy network, 5 for the random one, 4 when everyone has 4 contacts.
- Long-run circulation is about feedback loops that keep reinfecting each other, not only about average degree.

APPROXIMATION	TRACKS	GOOD FOR
degree-based mean field	share infected among people with k contacts	hubs, heavy tails
edge-based / message passing	whether each edge has transmitted yet	curves and final size on tree-like networks
full simulation	every person	anything, at a cost

THREE THRESHOLDS

A threshold depends on which object encodes who meets whom: the mean (well mixed), κ (random network, SIR), or ρ(A) (network, SIS).

Interventions on the graph

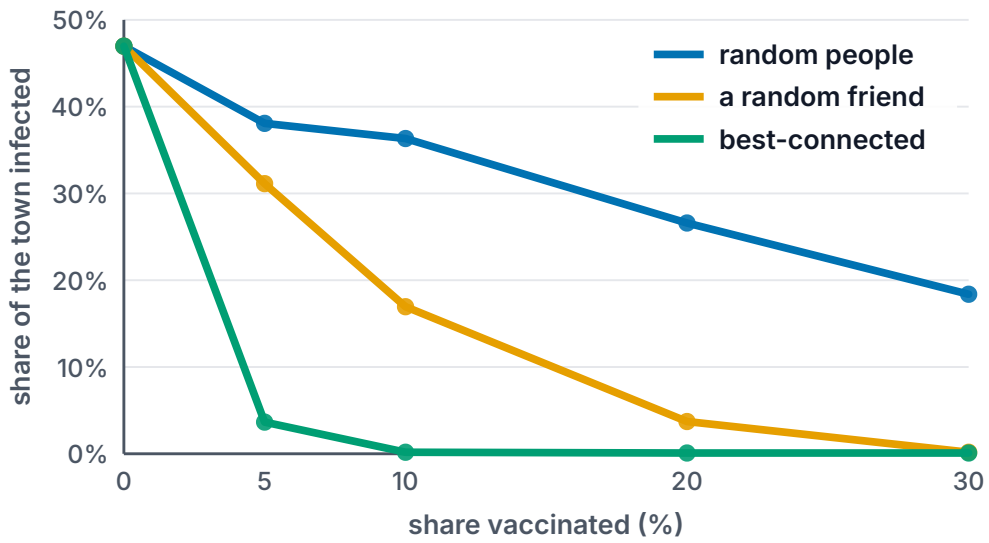

Our runs: hub-heavy network, T = 1/2, average of 200 runs per point.

- Vaccination removes nodes; distancing removes or weakens edges; tracing and quarantine follow and cut the edges of each case.
- With 20% vaccinated: at random, 27% are still infected; the best-connected 20%, 0.09%. But targeting needs the whole network.
- Acquaintance immunization (Cohen, Havlin and ben-Avraham, 2003): pick random people and vaccinate one of their friends. Friends are more often hubs. With 20%: 3.7%.

IN PLAIN WORDS

Networks pay off when contact structure matters: superspreading, households, targeted measures. Without contact data, assumptions do the work; often a hybrid (networks inside patches, travel between) is the compromise.

Check yourself: Part 2

CHECK YOURSELF

4. $\tau = 0.2$ per day along an edge and $\gamma = 0.1$ per day. What is T?

5. Using the threshold $T\,\kappa > 1$, would the outbreak in question 4 take off on a network where everyone has exactly 2 contacts? On one with 3?

6. Why does vaccinating a random friend of a random person work better than vaccinating a random person?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 3

Continuous space

Densities, diffusion and travelling waves.



PART 3 · CONTINUOUS SPACE

Reaction and diffusion

$$\frac{\partial S}{\partial t} = -\beta SI + D\nabla^2 S, \quad \frac{\partial I}{\partial t} = \beta SI - \gamma I + D\nabla^2 I, \quad \frac{\partial R}{\partial t} = \gamma I + D\nabla^2 R$$

- S, I, R become densities at each point x of a region. The reaction part is the local SIR; the diffusion part moves people.
- Diffusion is the large-scale limit of many small random steps. A walk with steps of length ℓ every Δt spreads with $D = \ell^2/(2\Delta t)$ in one dimension.
- Places can differ: β(x), γ(x) and D(x) can vary in space (cities, barriers, corridors).

IN PLAIN WORDS

Diffusion is a summary of many short, roughly random trips, not a literal physical process; D says how fast people mix with their neighbourhood.

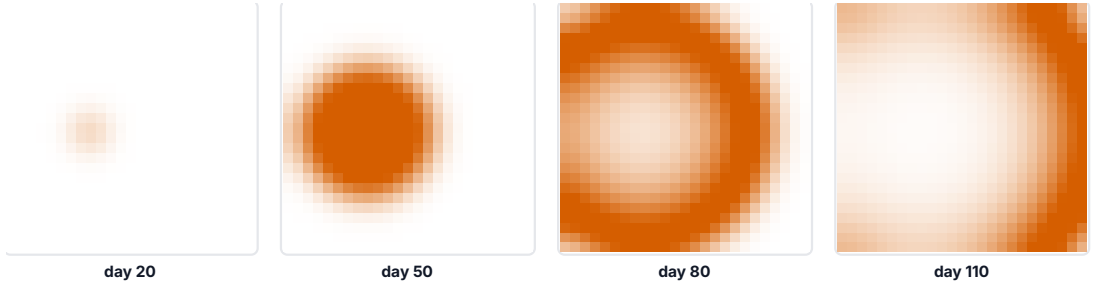

Our 2D run: one small seed in a 100 × 100 km square with closed borders; colour = share infected now. A ring moves outward: recovered behind, susceptible ahead.

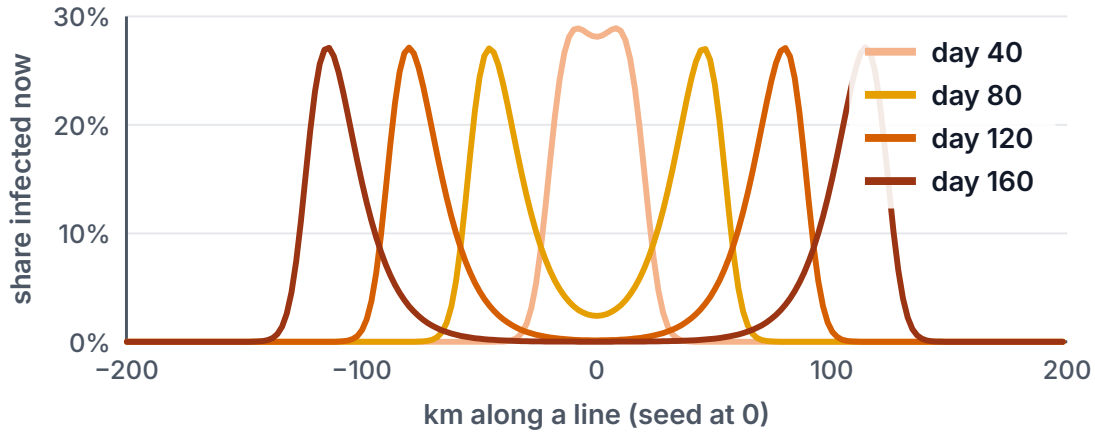

Our 1D run: the same model on a line; the wave keeps its shape as it travels.

The speed of the front

$$c^* = 2\sqrt{D\,r}, \qquad r = \beta - \gamma$$

- Fisher (1937) and Kolmogorov, Petrovsky and Piskunov (1937): a growing, diffusing quantity forms a travelling wave whose speed is set by movement and early growth.
- For SIR, linearize ahead of the front (S near 1): infections grow at r = β − γ and diffuse with D.
- Quadruple D or r: the speed doubles. Movement matters through a square root.

D (KM²/DAY)	R (PER DAY)	THEORY (KM/DAY)	MEASURED
1	0.2	0.89	0.87
4	0.2	1.79	1.75
1	0.4	1.26	1.23

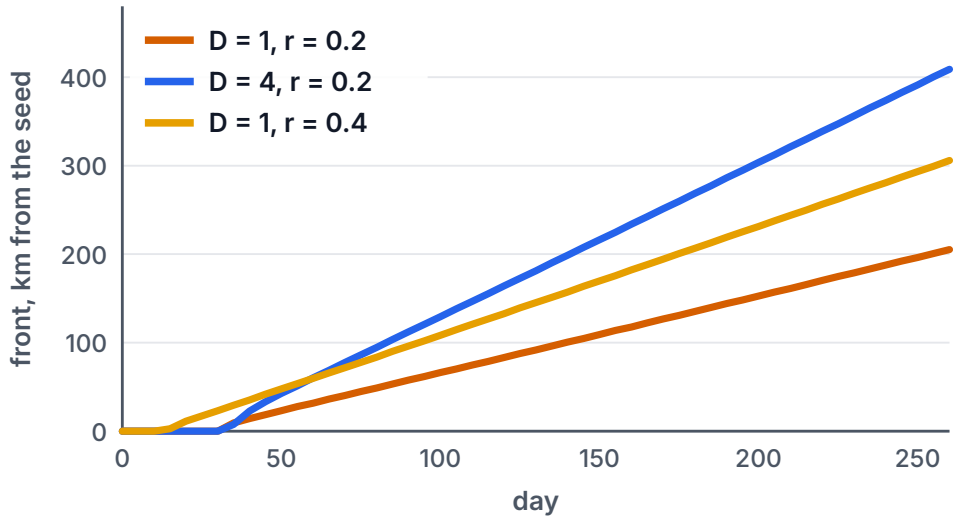

Our 1D runs: front = farthest point where half have been infected; speeds fitted over days 120–250. Fronts settle slowly, so measured speeds sit just below theory.

Borders are assumptions

- **Closed** (Neumann, no flux across the edge): an isolated region. Diffusion alone conserves the total.
- **Absorbing** (Dirichlet, density 0 at the edge): whoever reaches the edge leaves for good. A fixed positive value instead models a steady source of infection outside.
- **Leaky** (Robin): the outflow is proportional to the density at the edge, with an exchange rate κ: in between the two.
- A model with the wrong border can blame transmission for what was really travel, or miss constant reseeding.

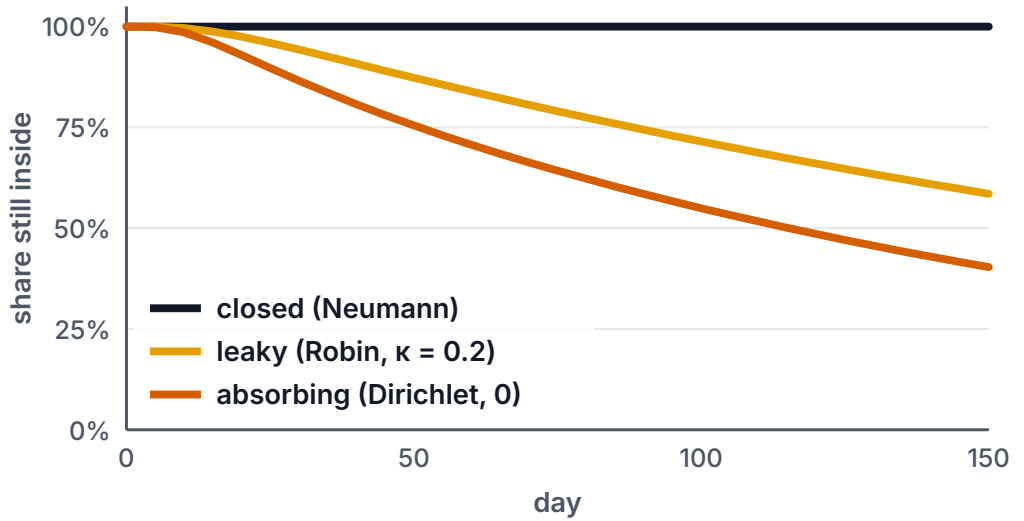

Our runs: the same bump of people diffusing (D = 1) in a 40 km region with each kind of border.

BORDER	STILL INSIDE AFTER 150 DAYS
closed	100.0%
leaky (κ = 0.2)	58.5%
absorbing	40.3%

PART 3 · CONTINUOUS SPACE

Drift and uneven landscapes

$$\frac{\partial I}{\partial t} = \beta S I - \gamma I - \nabla \cdot (\mathbf{v} I) + \nabla \cdot (D \nabla I)$$

- Directed movement along corridors adds an advection term. Our drift of 0.4 km/day: the front runs 1.31 km/day downstream and 0.51 upstream; theory c* ± v = 1.29 and 0.49.
- Spatially varying β, γ and D can slow, bend or pin fronts, and favourable pockets can keep infection going after the surroundings fade.

IN PRACTICE

Numerics: advection steepens fronts; schemes that respect the flow direction (upwinding) avoid fake oscillations.

STRIP WITH B = 0.05 (BELOW Γ)	DELAY OF THE WAVE
5 km	10 days
10 km	19 days
20 km	39 days
40 km	82 days

IN PLAIN WORDS

In the equations a low-transmission strip only delays the wave: tiny fractions of a person leak across and restart it. With whole people and chance, a wide gap can break the chain. Equations cannot tell a millionth of a person from zero.

PART 3 · CONTINUOUS SPACE

Long jumps

$$\frac{\partial I}{\partial t} = \beta S(x) \int K(x-y)\, I(y)\, dy - \gamma I$$

- Diffusion allows only small steps. Planes, trains and commuting add rare long jumps.
- Replace diffusion with a kernel K: how strongly an infected person at y affects people at x. Thin-tailed kernels behave like diffusion; fat-tailed ones seed new foci far ahead of the front.
- Our runs (same scale, 2 km): a Gaussian kernel gives a steady front, about 0.76 km/day; a fat tail ($\propto 1/d^4$) accelerates, from about 1.3 to 84 km/day.

IN PLAIN WORDS

Real spread is often hybrid: smooth local waves plus discontinuous seeding by travel between hubs. Discretize a kernel model and it becomes a weighted graph: Part 1 again.

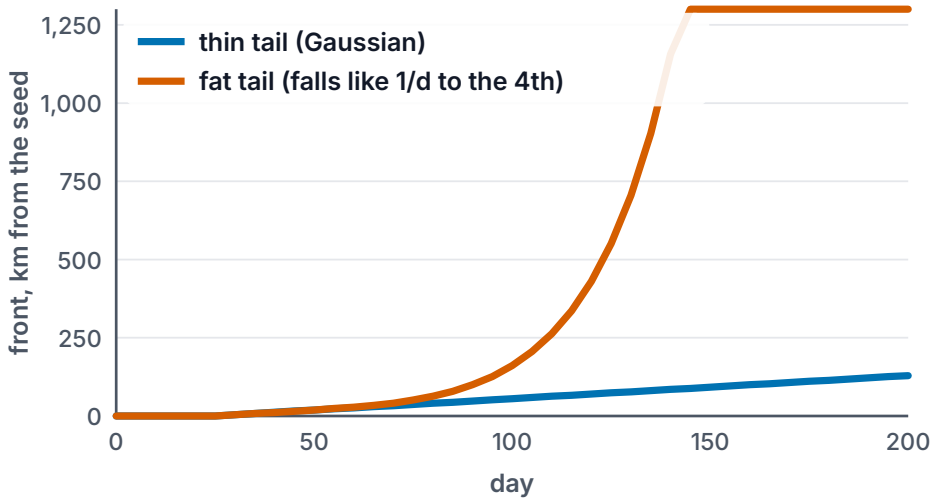

Our 1D runs: force of infection through a kernel; front = farthest point where half have been infected (fat tail capped at 1,300 km in the plot).

Check yourself: Part 3

CHECK YOURSELF

7. D = 2 km²/day, β = 0.5 and γ = 0.25 per day. Estimate the wave speed.

8. Which border condition fits an island with no ferries? A region next to a large city with a big ongoing outbreak?

9. Why can a fat-tailed kernel make the front speed up instead of settling to a constant speed?

TIP

Answers are at the end of the deck.

Applied Computational Mathematics Lab

PART 4

Norms and stability

Bounding growth without solving everything.



One generation at a time

$$x_{t+1} = A\,x_t, \qquad \text{grows} \iff \rho(A) > 1$$

- Near the start, almost everyone is susceptible, so new cases per district depend linearly on current cases in all districts.
- Column q of A: what one case in district q causes, district by district. This is the next-generation matrix of Module 3, with districts as types.
- ρ(A), the spectral radius, is the largest absolute eigenvalue: the long-run growth factor per generation.

NORMS

$\|x\|_1$ = total cases; $\|x\|_\infty$ = the worst district; $\|x\|_2$ = the Euclidean size. A matrix norm is the most A can amplify a vector, measured in that norm.

INDUCED NORM	FORMULA	MEANING HERE
$\|A\|_1$	largest column sum	most cases one case causes
$\|A\|_\infty$	largest row sum	most pressure one district receives
$\|A\|_2$	largest singular value	worst amplification of the Euclidean size

The guarantee

$$\rho(A) \le \|A\| \quad \text{for every induced norm}$$

	D1	D2	D3	D4	D5	ROW SUM
d1	0.55	0.20	0.00	0.00	0.10	0.85
d2	0.15	0.50	0.15	0.00	0.00	0.80
d3	0.00	0.20	0.45	0.20	0.00	0.85
d4	0.00	0.00	0.20	0.50	0.25	0.95
d5	0.25	0.00	0.00	0.25	0.60	1.10
column sum	0.95	0.90	0.80	0.95	0.95	

- $\|A\|_1$ = 0.95 < 1: decay is proven without an eigenvalue. $\|A\|_\infty$ = 1.10: that norm cannot decide. The truth: ρ(A) = 0.92 ($\|A\|_2$ = 0.93).

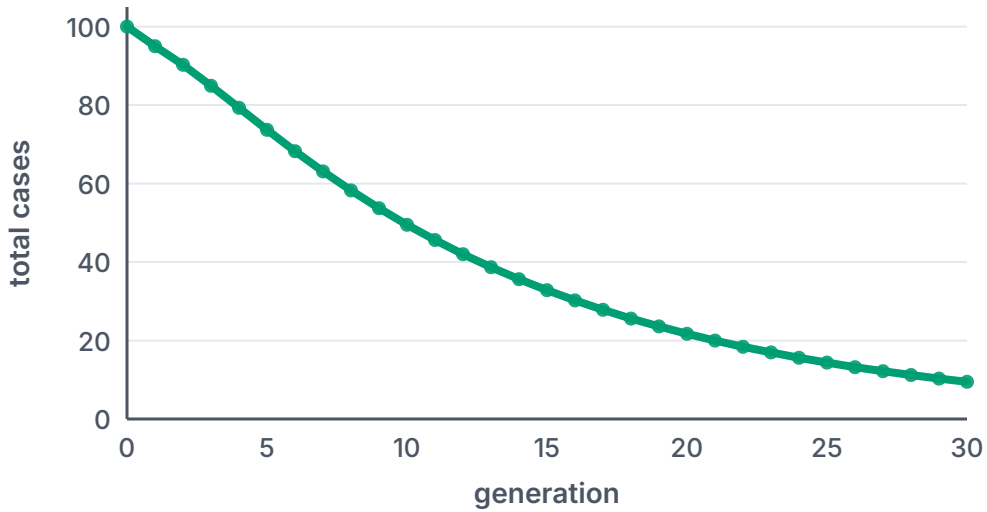

Our arithmetic: 100 cases in district 5; the total falls to 50 by generation 10.

IN PLAIN WORDS

A norm below one is a certificate. A norm above one proves nothing: try another norm, or compute ρ(A).

Short run, long run

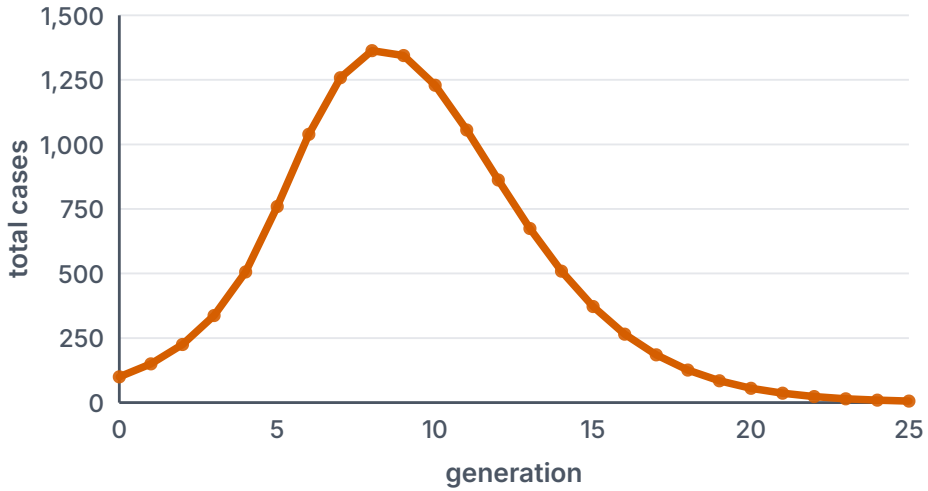

Our arithmetic: six districts on a route; each case causes 0.5 at home and 1.0 in the next district. 100 cases in the first district.

- ρ(A) = 0.5 (the matrix is triangular: the eigenvalues are the diagonal), so the outbreak dies out.
- But $\|A\|_1$ = 1.5: totals can grow by up to 50% per generation, and do: 1363 cases at generation 8 before the fall.
- Such non-normal matrices (strong one-way coupling) show transient growth. $\|x_t\| \leq \|A\|^t \|x_0\|$ bounds it step by step.
- Gelfand's formula: $\|A^k\|^{1/k} \to \rho(A)$ as k grows; slowly here.

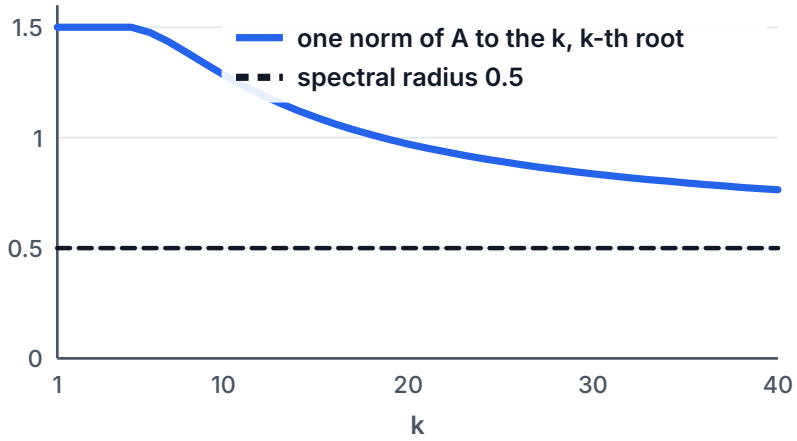

Keep and export: a simple rule

$$A = a\,I + b\,W, \quad W \text{ column-stochastic } \Rightarrow \ \|A\|_1 = \rho(A) = a + b$$

- Each case causes a at home and b elsewhere, split over the other districts by a travel matrix W whose columns sum to one.
- Every column of A sums to a + b, and for a non-negative matrix with equal column sums that sum is the spectral radius: the bound is exact.
- So local measures (lower a) and travel measures (lower b) trade one for one.

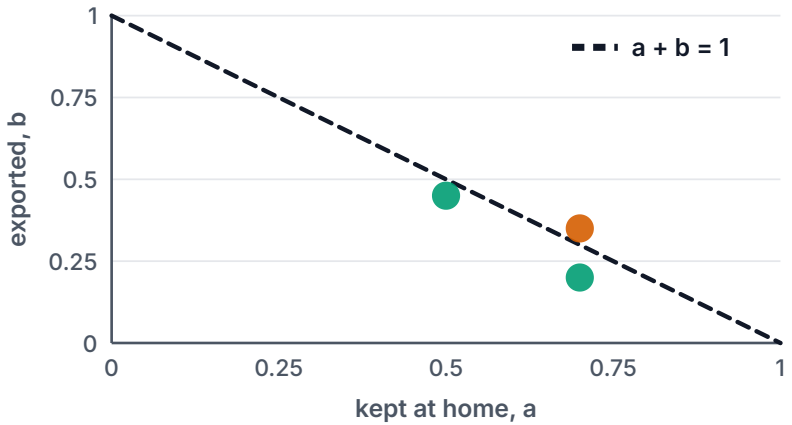

Below the dashed line every case causes fewer than one new case.

KEEP A	EXPORT B	P(A) = $\|A\|_1$	$\|A\|_\infty$
0.7	0.2	0.90	0.98
0.7	0.35	1.05	1.19
0.5	0.45	0.95	1.13

Our arithmetic with a 4-district travel matrix W; the row sums differ, so $\|A\|_\infty$ is not exact.

TRY IT

Local transmission is cut to a = 0.6. How large can the exported share b be while decay stays guaranteed?

Norms in continuous space

$$\frac{d}{dt}\|I\|_2^2 = 2\int(\beta S - \gamma)I^2 - 2D\int|\nabla I|^2 \;\leq\; 2r\|I\|_2^2$$

- With closed borders, diffusion only smooths: the L2 norm of the infected density cannot grow faster than r = β − γ.
- Our 1D run: r = 0.2; the fastest measured growth of $\|I\|_2$ was 0.17 per day.
- With absorbing borders, diffusion carries infection out of the region. On a strip of width L, an outbreak can grow only if $L > L_c = \pi\sqrt{D/r}$ (the critical patch size).

STRIP WIDTH	MEASURED GROWTH PER DAY	THEORY R − DΠ²/L²
0.8 L_c	-0.1125	-0.1125
1.2 L_c	0.0611	0.0611

Our runs: D = 1, r = 0.2, so L_c = 7.02 km; linearized model with absorbing edges.

IN PLAIN WORDS

Small enough habitats, or areas that lose people fast enough at their edges, cannot sustain an outbreak on their own.

From continuous space back to a graph

TURNING A CONTINUOUS MODEL INTO A PATCH MODEL

1. Partition the region into areas where people mix fast compared with epidemic growth.
2. Add up S, I and R inside each area: node states.
3. Turn flows across shared borders into travel rates (finite volumes: D / h^2 between neighbours of width h).
4. For kernels, average the kernel over pairs of areas: a mixing matrix.

TOWNS ON A 120 KM LINE	3	6	12	24	48	120
day the far third is half infected	75	94	114	128	136	139

Our runs: the continuous model reaches that point on day 140.

- Coarse graphs make the wave too fast: inside a node, infection spreads instantly.
- Fewer nodes are cheaper to run, calibrate and explain; check what you lose.

TRY IT

Why does each extra node in the graph slow the wave down, towards the continuous answer?

Check yourself: Part 4

CHECK YOURSELF

10. A = [[0.6, 0.3], [0.2, 0.5]] (columns: from district 1, from district 2). Compute $\|A\|_1$, $\|A\|_\infty$ and ρ(A). Does the outbreak die out?

11. ρ(A) = 0.7 but $\|A\|_1$ = 1.4. Can the total number of cases grow?

12. D = 0.5 km²/day and r = 0.125 per day. What is the critical width of a strip with absorbing edges?

TIP

Answers are on the next slides.

ANSWERS

Answers: warm-up and Parts 1–2

- **(a)** False: travel changes when cities are hit, not how many are infected. **(b)** False: the threshold depends on the excess degree (and clustering), not just the mean. **(c)** True: the speed is $2\sqrt{Dr}$. **(d)** False: non-normal matrices can grow for a while (Part 4).
- **1.** When its wave comes (earlier); the final size stays about the same.
- **2.** Repeated sparks from a linked city with an outbreak; or visitors bringing infection pressure (mixing).
- **3.** (2 + 4 × 0.5) / 5 = 0.8: below one, so it dies out (assuming equal sizes and full mixing).
- **4.** 0.2 / 0.3 = 2/3. **5.** Everyone 2: κ = 1, T κ = 0.67, no. Everyone 3: κ = 2, T κ = 1.33, yes.
- **6.** Friends are reached by following an edge, so they tend to have more contacts: the trick finds hubs without knowing the network.

TRY-IT ANSWERS

Try-it answers: 50,000 × 0.02 = 1,000 people a day, half to each other patch. Full mixing: every row of M is (1/P, …, 1/P), so diag(R)M has rank one and ρ = the mean of R. Half 2, half 6: mean 4, $\langle k^2 \rangle = 20$, κ = 4, $T_c = 0.25$ (vs 1/3).

Answers: Parts 3–4

- **7.** r = 0.25, c* = $2\sqrt{2 \times 0.25}$ = 1.41 km/day.
- **8.** Island: closed (no flux). Next to a big outbreak: a fixed positive value at the border (constant source), or a leaky border with outside infection.
- **9.** Rare long jumps seed new foci ahead of the front, and each new focus grows and seeds further; with fat tails the farthest seeds dominate, so the front keeps accelerating.
- **10.** Column sums 0.8 and 0.8: $\|A\|_1$ = 0.8. Row sums 0.9 and 0.7: $\|A\|_\infty$ = 0.9. Eigenvalues 0.8 and 0.3: ρ(A) = 0.8. It dies out (proven by either norm).
- **11.** Yes, for a while: up to 40% per generation, as in the route example; in the long run it decays.
- **12.** $\pi\sqrt{0.5/0.125} = 2\pi \approx 6.3$ km.

TRY-IT ANSWERS

Try-it answers: b < 0.4. More nodes: each node mixes instantly inside, which acts like extra fast diffusion over its width; shrinking the width removes that artificial head start.

CREDITS

Sources and credits

- This deck accompanies Chapter 4 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026), the course textbook. It follows the chapter's topics in our own words; no text, figures, tables, exercises or code are copied from the book.
- The text, figures, examples and simulations were written, drawn and run for ACML; nothing is copied from the book. Numbers labelled "our run" come from our own code; counts from agencies and papers are credited where they appear.
- All numbers, curves and toy calculations are our own runs (production/explains/sims/epi_m04_sims.py), chosen to show mechanisms, not to estimate any real disease.
- Standard results of the field described in our own words: effective distance (Brockmann and Helbing, Science, 2013), network thresholds and percolation (Newman, 2002), acquaintance immunization (Cohen, Havlin and ben-Avraham, 2003), travelling-wave speed (Fisher; Kolmogorov, Petrovsky and Piskunov, 1937), the critical patch size (Kierstead, Slobodkin and Skellam).
- Companion videos: **Epidemic Modeling**, module 4, episodes 13–16. More from the lab: acml.teddylazebnik.com

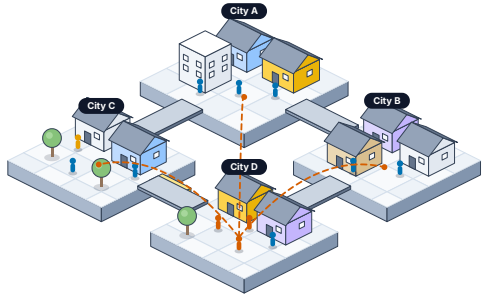

Applied Computational Mathematics Lab

END OF MODULE 4

Spatio-temporal models

YOU COVERED

- Cities linked by travel: metapopulation models
- Networks of people
- Epidemics in continuous space
- Bounding the spread: norms and stability

Module 4 complete

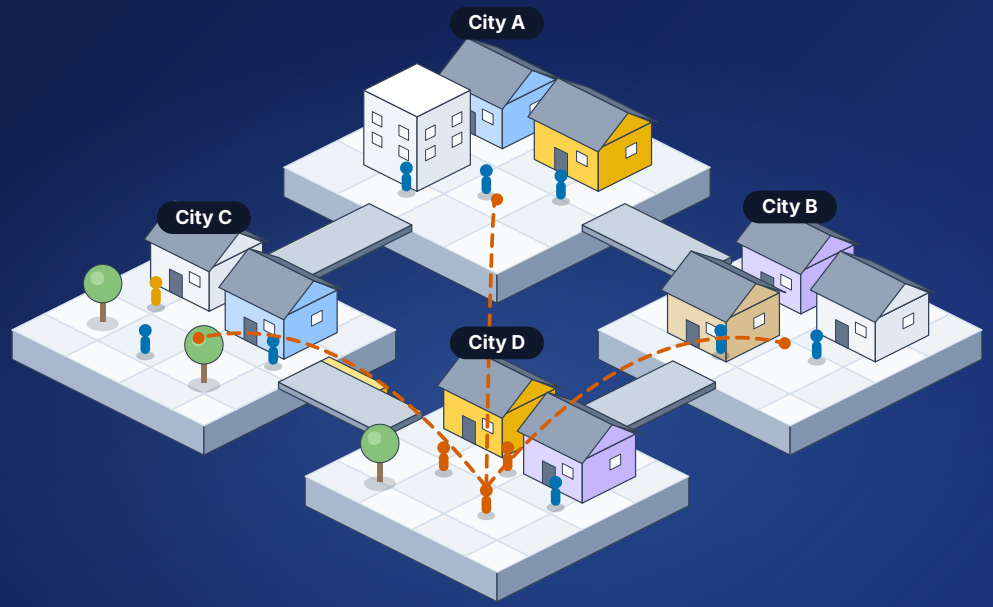

UP NEXT · MODULE 5

Multi-strain and evolving pandemics

Strains, reinfection and mutation

Based on Chapter 4 of T. Lazebnik, Epidemic Modeling: Principles, Algorithms, and Applications (Academic Press, 2026)