

1 · Why Traffic Is Weird

A jam out of nothing · trajectories · flow, density and speed

Part 1 · A jam out of nothing


The puzzle

In 2008, researchers in Japan asked 22 drivers to drive around a 230 m circle, all at the same steady speed. There was no bottleneck, no accident and no traffic light. (Sugiyama et al., *New Journal of Physics*, 2008)

What happened next?

i Before you watch: predict

Will a jam form? If one does, which way will it move: **with** the traffic, or **against** it? Make a guess before you go on.



Watch it happen

We can recreate the experiment below. Each dot is a car, coloured by its speed. The car with the black ring is **you**. Just wait and watch.

Interactive simulation: see the web version of these notes. Cars on a ring road start at equal spacing and speed. Tiny fluctuations in driving grow into a stop-and-go wave that travels backward around the ring.

Follow one driver

Keep watching the ring on the previous slide. This is your speed, the car with the black ring. Nobody ahead of you ever had a reason to stop, yet you stop and go, again and again.

Interactive figure: see the web version of these notes. The black car stops and starts again and again, although nobody ahead of it had a reason to stop.

The trace is live: it keeps updating while the ring runs. Stopping and starting again and again, with no bottleneck anywhere, is what a phantom jam feels like from the driver's seat.


What just happened?

- A jam formed within minutes, exactly as in the 2008 experiment. **Nobody caused it.** It grew out of **tiny fluctuations** that each driver over-corrected.
- The red band in the right-hand plot slopes **downward**: the jam moves backward while every car moves forward.
- **Try it:** raise *driver responsiveness* to 1.5. The jam dissolves. Automated vehicles with well-designed cruise control could do exactly this (Lecture 9).

The plot on the right is a **time-space diagram**: time runs left to right and position runs bottom to top. Each car leaves a trail of dots. When the trails bunch up and turn red, cars are stopped in a jam. The jam's boundary tilts downward because cars leave the front of the jam while new cars join its back. The jam moves backward at a speed of its own, whatever the individual cars do. By Lecture 5 you will be able to compute that speed with pencil and paper.

It happens on real freeways, every morning

This is not just a simulation. In 2005 the Federal Highway Administration filmed US-101 in Los Angeles and recorded the position of every vehicle, ten times a second. Here are four minutes of one lane:



The red band is a stop-and-go wave, just like the one on the ring. It travels backward, against the traffic, at roughly 15–20 km/h. The same speed shows up on freeways all over the world. **In Lab 1 you will load these data yourself and measure the wave's speed.**

The 1959 experiment: one number decides

In 1959, Robert Herman, Elliott Montroll, Renfrey Potts and Richard Rothery at General Motors Research wrote down the simplest possible driver. Each driver adjusts speed to match the car ahead, but only after a **reaction time** T :

$$\ddot{x}_{n+1}(t) = \lambda [\dot{x}_n(t - T) - \dot{x}_{n+1}(t - T)].$$

Here λ is how strongly the driver responds. Everything depends on one number, $C = \lambda T$: **how strongly the driver reacts** $\times$ **how late**.

Below, the leader of a line of 8 cars brakes briefly, then speeds up again. Watch what happens to the cars behind as you change C .

Try it: C decides

Interactive figure: see the web version of these notes. Herman et al. found three thresholds. For $C \leq 1/e \approx 0.37$ a follower returns smoothly to speed. For $C < \pi/2$ each follower is stable on its own. But only for $C < 1/2$ does a disturbance fade as it passes down a line of cars. Between $1/2$ and $\pi/2$ it grows from car to car: a phantom jam.

Stable driver, unstable line

The key idea: stability of one car is not stability of the line. Between $C = 1/2$ and $C = \pi/2$, every driver on their own settles down perfectly well, and yet the line of cars amplifies disturbances. That is exactly what happened on the ring road. It is also why the ring-road jam dissolved when you made drivers more attentive: faster reactions mean a smaller T , and so a smaller C . You will derive all three thresholds yourself in Lecture 9.

In Lab 2 you will rebuild this 1959 experiment from scratch, find the thresholds, and discover when a line of perfectly reasonable drivers crashes.

Your turn: be the smart car

The jam is back. This time **you** drive the black car. Tick the box, then choose one cruising speed; your car holds it smoothly, never closer to the car ahead than is safe. Can one driver dissolve the jam for all 22? Watch the jam meter.

Interactive game: see the web version of these notes. One car that drives smoothly at a modest, steady speed can dissolve the ring-road jam, and everyone's average speed rises.


The winning strategy is to cruise at about 2–3 m/s. That is slower than most students choose at first. The smart car opens a gap in front of it and uses that gap to absorb each wave instead of braking. The jam meter falls from about 2.5 to 0.3, and the average speed of all 22 cars goes *up*, from about 7.5 to 9–10 km/h. This is the idea behind the 2017 field experiment in Tucson (Stern et al., *Transportation Research Part C*, 2018), where a single automated vehicle damped stop-and-go waves on a ring of about 20 cars. In Lab 12 you put one reinforcement-learning car among rule-following cars and see whether it finds a strategy like this on its own.

Part 2 · Reading traffic: trajectories

A trajectory is a story

A vehicle's **trajectory** is its position over time, $x(t)$. Draw it with time across and position up, and its shape tells the story of the trip.

What you see	What it means
Slope of the line	speed, $v = dx/dt$. Steeper is faster
Flat (horizontal) line	stopped
Bending flatter	braking
Bending steeper	speeding up



Students often want position on the horizontal axis because that is how a map looks. Insist on the convention: time runs across, position runs up. Every result in the macroscopic part of the course (waves, shocks, queues) is read off this picture.

Try it: drive one car


A car cruises, brakes for a red light at 200 m, waits, then drives off. Change how fast it goes and how long it waits, and watch the story change shape.

Interactive figure: see the web version of these notes. A car's trajectory is straight while it cruises, bends flatter as it brakes, is flat while it waits, and bends steeper as it speeds up.

Many trajectories: the time-space diagram

Draw every vehicle on one picture and you get a **time-space diagram**.

- **Spacing** is the *vertical* gap between two trajectories: how far apart two cars are at one instant.
- **Headway** is the *horizontal* gap: how much time passes between two cars at one point.
- In a single lane, trajectories **never cross**. A crossing means someone changed lanes, or crashed.



Spacing (m) and headway (s) are the microscopic versions of density and flow. Average spacing is $1/k$, and average headway is $1/q$. The next slide shows both on a live diagram.

Live: a red light and eight cars

Eight cars arrive at a traffic light, which turns red at 16 s and green at 46 s. Press **Play**, or drag the time slider yourself. The road is on the left; its time-space diagram is on the right. Colour shows speed: red is slow, green is fast.

Interactive animation: see the web version of these notes. Eight cars approach a signal. During the red phase their trajectories turn flat one after another, and after the green they bend upward again one after another.

Point out the two lines drawn on the diagram. The loop detector is a horizontal line: one position, all times. The drone snapshot is a vertical line: one instant, all positions. A GPS probe would be a single trajectory. Every sensor sees the diagram through its own slice, and Lecture 2 builds on exactly this.

Can you read it?

Use the red-light diagram on the previous slide.

1. While the light is red, the points where trajectories turn flat move **down** the diagram, one car after another. What is moving down, and roughly how fast?
2. When the light turns green, the cars do not all start together. Where does the “start moving” signal travel?
3. The stopped cars have the smallest spacing of all, yet during the red no car passes the stop line. So is a queue “high density, zero flow”? Where is the flow highest just after the green?

💡 Answers

1. The **back of the queue**. It moves upstream, against the traffic, as each arriving car stops a little further back. Its speed is the slope of the line joining the points where the trajectories turn flat. This is a **shock wave**, the subject of Lecture 5.
2. **Backward**, from the stop line upstream. Each car can start only after the car ahead has moved. This start-up wave travels upstream at roughly 15–20 km/h, like the waves on the ring road and on US-101.
3. Yes: the queue is at jam density with zero flow. Just after the green, flow is highest **at the stop line**, where cars leave the queue at the road's capacity. That is exactly the point where a traffic engineer would measure saturation flow.

Trajectories at work: traffic signals

The red-light diagram you just read is exactly how engineers now evaluate real signals, using trajectories from connected vehicles and phones instead of detectors in the pavement.

- **Retiming signals with a few percent of cars.** Henry Liu's group at Michigan retimed 34 signals in Birmingham, Michigan using only GM vehicle trajectories, about 7% of traffic. Delay fell by up to about 20% and stops by up to about 30% (Wang et al., *Nature Communications*, 2024). Earlier, with DiDi, the same group optimized signals in Jinan, China from ride-hailing trajectories (Zheng et al., TRB, 2018).
- **Signal report cards.** Purdue's group (Darcy Bullock) computes arrivals on green, split failures and level of service from commercial connected-vehicle data, with no detectors at all (Saldivar-Carranza et al., *TRR*, 2021), including diamond and diverging diamond interchanges (*J. Transportation Technologies*, 2021 and 2022).
- **Coordinating an arterial.** Jianyuan Xu and Zong Tian at the University of Nevada, Reno use trajectories to evaluate arterial coordination (Xu et al., *IJTST*, 2025) and to split a corridor into coordinated groups by where drivers actually go (Xu and Tian, *TRR*, 2023).

More in Lecture 5 (queues and shock waves) **and Lecture 12** (a real signal controller in the loop).

The key idea: no single connected vehicle tells you much, but thousands of them, overlaid on one time-space diagram for the same intersection and the same time of day, reveal the queue, the arrivals on green and the delay. That is the red-light diagram on the previous slides, filled in with real data. Purdue calls its version the "Purdue Probe Diagram".

Trajectories at work: automated vehicles

- **Learning how people drive.** Car-following and lane-changing models are calibrated on trajectories. Deep learning can go further: Zhu, Wang and Wang (*TR-C*, 2018) trained a human-like car-following model on 2,000 events from the Shanghai Naturalistic Driving Study.
- **Testing automated vehicles.** Dangerous situations are rare in real driving, so testing on public roads alone would take enormous mileage. Feng et al. (*Nature*, 2023) trained background vehicles on naturalistic driving data, then edited their behavior to keep only the safety-critical moments, making tests far faster while staying statistically unbiased.
- **Naturalistic driving data** follow *one driver* for months, instead of *every vehicle* on one stretch of road. The largest, the SHRP2 study (2010–2013), instrumented the cars of more than 3,400 US drivers: about 5.4 million trips and over 1,500 crashes on video.

More in Lecture 8 (car-following), **Lecture 9** (automated vehicles and phantom jams), **Lecture 10** (calibration) **and Lecture 12** (reinforcement learning).

Trajectories beyond this course: travel demand

Can trajectories feed **travel demand models**? Yes, but not the ones in Lab 1.

	High-frequency trajectories	Sparse GPS traces
Examples	NGSIM, highD	phones, ride-hailing, connected vehicles
Sampling	10–25 times a second	every few seconds to minutes
Coverage	every vehicle, a few hundred metres	some vehicles, the whole city, whole trips
Good for	traffic flow, driver behavior	origins and destinations, route choice

Sparse traces are now a major input to demand models: origin–destination matrices from phone location data (Çolak et al., *TRR*, 2015) and route choice from smartphone GPS (Bierlaire, Chen and Newman, *TR-C*, 2013). The catch: only *some* travelers are observed, and not a random sample, so the data must be expanded and corrected for bias.

Where do trajectories come from?

Dataset	Where	How	What
NGSIM (FHWA, 2005–06)	US-101 and Lankershim Blvd (Los Angeles), I-80 (Emeryville, CA), Peachtree St (Atlanta)	video from tall buildings, 10 Hz	every vehicle, about 500–700 m
highD (RWTH Aachen, 2017–18)	German highways near Cologne	drone video, 25 Hz	110,000 vehicles, 16.5 h
inD, roundD, exiD	German intersections, roundabouts, ramps	drone video	cars, trucks, cyclists, pedestrians
pNEUMA (EPFL, 2018)	central Athens	a swarm of 10 drones	almost half a million trajectories
I-24 MOTION (Vanderbilt)	4.2 miles of I-24, Nashville	276 pole-mounted cameras	every vehicle, every morning
TGSIM (FHWA)	Chicago, Washington DC	helicopters and fixed cameras	human drivers next to automated vehicles
Waymo Open Motion, Argoverse 2	US cities	sensors on self-driving cars	the scene around one vehicle

Full trajectories of every vehicle are rare and precious. Most traffic data are only **slices** of the time-space diagram, such as a line at one point or a snapshot at one instant. What each slice can and cannot tell us is **Lecture 2**. In **Lab 1** you will load real NGSIM trajectories yourself.

NGSIM: why the data had to be rebuilt

NGSIM turned video into trajectories automatically. Positions were close, but **speeds and accelerations come from differentiating positions**, and differentiating amplifies every small error.

- Raw NGSIM contains large measurement errors (Punzo, Borzacchiello and Ciuffo, *TR-C*, 2011): accelerations no car can produce, speeds that jump in steps, and followers that drive *through* their leader (Coifman and Li, *TR-B*, 2017).
- **Smoothing**: filter the positions before differentiating (Thiemann, Treiber and Kesting, *TRR*, 2008).
- **Reconstruction**: Montanino and Punzo (*TR-B*, 2015) rebuilt the I-80 data so every trajectory obeys physics and every leader–follower pair stays consistent. With the rebuilt data, models reproduced the observed jams better.
- **Re-extraction**: Coifman and Li (*TR-B*, 2017) went back to the original video and re-tracked vehicles by hand. They argued that some errors cannot be cleaned away, and they found 11% more vehicles.

The lesson: calibrate a car-following model on raw NGSIM and it learns the *noise*, not the driver. Always check a trajectory's speed and acceleration before trusting it. **You will see this in Lab 1, and it matters again in Lectures 8 and 10.**

Ask students: if position is accurate to half a metre, and samples are 0.1 s apart, how large an error in speed can one bad sample create? Half a metre over 0.1 s is 5 m/s, or 18 km/h. Differentiate again for acceleration and the error is 50 m/s², five times gravity. That is why raw accelerations look absurd.

Further reading: trajectories

- Bierlaire, M., Chen, J. and Newman, J. (2013). A probabilistic map matching method for smartphone GPS data. *Transportation Research Part C* 26, 78–98. <https://doi.org/10.1016/j.trc.2012.08.001>
- Çolak, S., Alexander, L. P., Alvim, B. G., Mehndiratta, S. R. and González, M. C. (2015). Analyzing cell phone location data for urban travel. *Transportation Research Record* 2526, 126–135. <https://doi.org/10.3141/2526-14>
- Coifman, B. and Li, L. (2017). A critical evaluation of the Next Generation Simulation (NGSIM) vehicle trajectory dataset. *Transportation Research Part B* 105, 362–377. <https://doi.org/10.1016/j.trb.2017.09.018>
- Feng, S., Sun, H., Yan, X., Zhu, H., Zou, Z., Shen, S. and Liu, H. X. (2023). Dense reinforcement learning for safety validation of autonomous vehicles. *Nature* 615, 620–627. <https://doi.org/10.1038/s41586-023-05732-2>
- Krajewski, R., Bock, J., Kloeker, L. and Eckstein, L. (2018). The highD dataset. *IEEE ITSC*, 2118–2125. <https://doi.org/10.1109/ITSC.2018.8569552>
- Montanino, M. and Punzo, V. (2015). Trajectory data reconstruction and simulation-based validation against macroscopic traffic patterns. *Transportation Research Part B* 80, 82–106. <https://doi.org/10.1016/j.trb.2015.06.010>
- Punzo, V., Borzacchiello, M. T. and Ciuffo, B. (2011). On the assessment of vehicle trajectory data accuracy and application to the NGSIM program data. *Transportation Research Part C* 19(6), 1243–1262. <https://doi.org/10.1016/j.trc.2010.12.007>
- Saldivar-Carranza, E., Li, H., Mathew, J., Hunter, M., Sturdevant, J. and Bullock, D. M. (2021). Deriving operational traffic signal performance measures from vehicle trajectory data. *Transportation Research Record* 2675(9), 1250–1264. <https://doi.org/10.1177/03611981211006725>
- Saldivar-Carranza, E., Li, H. and Bullock, D. M. (2021). Diverging diamond interchange performance measures using connected vehicle data. *Journal of Transportation Technologies* 11, 628–643. <https://doi.org/10.4236/jtts.2021.114039>
- Saldivar-Carranza, E., Rogers, S., Li, H. and Bullock, D. M. (2022). Diamond interchange performance measures using connected vehicle data. *Journal of Transportation Technologies* 12, 475–497. <https://doi.org/10.4236/jtts.2022.123029>
- Thiemann, C., Treiber, M. and Kesting, A. (2008). Estimating acceleration and lane-changing dynamics from Next Generation Simulation trajectory data. *Transportation Research Record* 2088, 90–101. <https://doi.org/10.3141/2088-10>
- Wang, X., Jerome, Z., Wang, Z., Zhang, C., Shen, S., Kumar, V. V., Bai, F., Krajewski, P., Deneau, D., Jawad, A., Jones, R., Piotrowicz, G. and Liu, H. X. (2024). Traffic light optimization with low penetration rate vehicle trajectory data. *Nature Communications* 15, 1306. <https://doi.org/10.1038/s41467-024-45427-4>
- Xu, J. and Tian, Z. (2023). OD-based partition technique to improve arterial signal coordination using connected vehicle data. *Transportation Research Record* 2677, 252–265. <https://doi.org/10.1177/03611981221098692>
- Xu, J., Tian, Z., Wang, A., Xie, G. and Valenzuela, L. (2025). Development and assessment of trajectory-based arterial through percent arrivals on red for arterial signal coordination performance evaluation. *International Journal of Transportation Science and Technology* 18, 131–147. <https://doi.org/10.1016/j.ijtst.2024.06.001>
- Zheng, J., Sun, W., Huang, S., Shen, S., Yu, C., Zhu, J., Liu, B. and Liu, H. X. (2018). Traffic signal optimization using crowdsourced vehicle trajectory data. *TRB 97th Annual Meeting*, paper 18-05789. <https://trid.trb.org/View/1496953>


- Zhu, M., Wang, X. and Wang, Y. (2018). Human-like autonomous car-following model with deep reinforcement learning. *Transportation Research Part C* 97, 348–368. <https://doi.org/10.1016/j.trc.2018.10.024>

Part 3 · Three numbers describe traffic

Flow, density and speed

Variable	Symbol	Units	Measures
Flow	q	veh/h	vehicles passing a point per hour
Density	k	veh/km	vehicles on a stretch of road
Speed	v	km/h	how fast they move

Flow q



a line across the road,
watched for an hour $\rightarrow$ veh/h

Density k



a stretch of road,
photographed once $\rightarrow$ veh/km

In a steady stream they are linked by

$$q = k v.$$

Check it on the ring: 22 cars on 230 m is a density of about 96 veh/km. If they all drive at 6 m/s (21.6 km/h), the flow past any point is $q = 96 \times 21.6 \approx 2070$ veh/h.

Try it: the fundamental diagram

Greenshields (1935) proposed that speed falls linearly as density rises: $v = v_f (1 - k/k_j)$. Then $q = v_f k (1 - k/k_j)$. Move the sliders and watch the road's capacity change.

Interactive figure: see the web version of these notes. With Greenshields' model, capacity is $q_{\max} = v_f k_j / 4$, reached at the critical density $k_c = k_j / 2$.

Part 4 · Wrap-up

This semester, you will learn to...

- **Explain** where jams come from, how fast they move, and when they clear (Lectures 3–6)
- **Build** your own traffic simulator in Python, piece by piece (all semester)
- **Measure** traffic from real data: drone and roadside video, freeway sensors (Lectures 1–4)
- **Fit** models to data, and know how much to trust them (Lecture 10)

- **Use AI** well: estimating and forecasting traffic, reinforcement learning, and knowing when a simple rule wins (Lectures 9, 11 and 12)
- **Understand** digital twins: simulations kept in sync with real roads, with real signal controllers plugged in (Lecture 12)

Three challenges

Use the simulations on this page. First to share the answer with the class, with a one-line explanation, gets bragging rights.

1. **Fewest cars.** With driver responsiveness at 0.5, what is the smallest number of cars on the ring that still produces a jam? Why does a jam need enough cars?
2. **Just attentive enough.** With 22 cars, what is the lowest driver responsiveness that keeps traffic smooth?
3. **The tipping point.** In the 1959 experiment, find the value of C at which the last car's disturbance is exactly as large as the first follower's. Theory says $1/2$. Do you get $1/2$? (Lab 2 explains why not.)

Summary

- Jams can appear out of nothing, and they travel backward.
- A driver who is stable alone can still be part of an unstable line of cars: $C = \lambda T$ decides.
- Trajectories and the time-space diagram contain all the information about a traffic stream.
- Flow, density and speed summarize it, and $q = kv$ links them.
- **Next week:** what sensors really measure, and why “average speed” has two different answers.

Notation

Symbol	Meaning	Units
t	time	s
$x(t)$	position of a vehicle along the road: its trajectory	m
$v = dx/dt$	speed	m/s or km/h
$a = d^2x/dt^2$	acceleration	m/s ²
s, h	spacing (distance to the car ahead), headway (time to the car ahead)	m, s
q	flow: vehicles passing a point per unit time	veh/h
k	density: vehicles per unit length of road	veh/km
v_f, k_j	free-flow speed, jam density	km/h, veh/km
$k_c, q_{\max}$	critical density, capacity (the top of the fundamental diagram)	veh/km, veh/h
T	driver reaction time (Herman et al.)	s
λ	driver sensitivity: how strongly a driver responds	1/s
$C = \lambda T$	the one number that decides stability	—