

2 · What Do Sensors Really Measure?

The same traffic, two different speeds · Edie's definitions · why your phone and the road disagree


Part 1 · How we watch traffic

Two ways to watch a road

Every traffic dataset starts with a sensor, and every sensor watches the road in its own way. Meet the two we compare today.

- **The loop detector** is a few turns of wire laid in a square cut in the pavement, about 1.8 m on a side, and wired to a cabinet by the road.
- **Why a loop?** Current flowing around a loop of wire makes a magnetic field above the lane: the loop is an *inductor*. When a car's metal body enters that field, the loop's inductance drops, and the electronics in the cabinet register a vehicle.
- So a loop sees **one point on the road, over time**: every car that crosses it.
- **The drone** films the road from above. It sees **the whole stretch, at one instant**: every car on it.

loop, seen from above:
3 turns of wire, about 1.8 m × 1.8 m




Loops are the workhorse of freeway monitoring: tens of thousands feed California's PeMS system alone. The next slide looks inside a loop; the rest of Part 1 surveys every other kind of sensor.

How a loop detector works


A closer look at the loop: how a coil of wire becomes counts and occupancy.

- A **coil of wire** in the pavement, wired to a roadside cabinet.
- A vehicle's metal body **lowers the loop's inductance**: the loop is **occupied**.
- Raw output: **on/off presence**. Every 20–30 s it becomes a **count** and an **occupancy** (share of time on).
- A **single loop cannot measure speed**; it is estimated (next slide).
- A **dual loop**, two loops a few metres apart, measures speed directly.

Seen from above: a wire loop cut into the lane, wired to the roadside cabinet



What the loop reports: count = number of pulses (5) occupancy = total on-time ÷ period



The electronics drive an alternating current through the coil; the drop in inductance shows up as a change in the loop's oscillation frequency. Loops have been the backbone of freeway monitoring for decades. California's PeMS system collects data from tens of thousands of them, and so does almost every freeway traffic management centre. They are cheap to run but expensive to repair, because the road must be cut open. That is one reason for the non-intrusive sensors in the field guide.


Classifying traffic sensors

By where they sit

- **In the pavement** (*intrusive*): loops, magnetometers, tubes, weigh-in-motion. Accurate, but the road must be closed to install or repair them.
- **Above or beside the road** (*non-intrusive*): radar, video, LiDAR, infrared, acoustic. Installed on poles and gantries, without touching the pavement.
- **In the vehicle**: phones, fleet GPS, connected vehicles.
- **In the air**: aerial photographs from aircraft, and drones.

By what they see

- **A point**: counts, occupancy, spot speeds.
- **Two points**: the same vehicle recognised twice, giving travel time.
- **An area**: every vehicle on a stretch, at an instant.
- **A moving vehicle**: one trajectory, or pieces of it.



Ask the class to place each sensor on the red-light diagram from Lecture 1: a horizontal line (point sensors), two horizontal lines (re-identification), a vertical line (a snapshot), or a trajectory (a probe). Only drones and aerial video recover whole trajectories, which is why datasets such as NGSIM and highD are so valuable.

A catalogue of traffic sensors

In the pavement

Sensor	How it works	Data	Pros	Cons
Inductive loop	wire coil in the pavement; a vehicle's metal lowers its inductance	count, occupancy; speed with two loops	mature, cheap to run, accurate counts	pavement must be cut; lane closures to repair; one loop cannot measure speed
Magnetometer	small puck senses the vehicle's disturbance of Earth's magnetic field	count, occupancy, some speed	quick to install in a small hole; often wireless	battery life; less reliable vehicle classes
Pneumatic tube	rubber tube across the road; each axle sends an air pulse	axle counts, vehicle class, speed with two tubes	cheap and portable, ideal for temporary counts	wears out; poor in slow or stop-and-go traffic; risky to install
Piezoelectric / weigh-in-motion	sensors in the pavement measure the force of each axle	axle weights, vehicle class, speed	the only way to weigh trucks at speed	expensive; needs careful calibration; pavement must be cut

Above or beside the road

Sensor	How it works	Data	Pros	Cons
Microwave radar	<i>Doppler</i> : frequency shift from moving vehicles. <i>FMCW side-fire</i> : ranges vehicles across the lanes	speed (Doppler); count, occupancy and speed by lane (FMCW)	all weather; one unit covers many lanes	Doppler cannot see stopped vehicles; tall trucks hide cars behind them
Video detection	cameras with image processing, now often deep learning	count, speed, class, queue length, incidents	rich information; one camera, many uses	glare, darkness, rain and snow; occlusion; privacy
LiDAR	laser pulses measure distance to every point in view	precise vehicle position, shape and class	very accurate 3-D picture	costly; affected by heavy rain and fog
Infrared	active: reflected infrared beam; passive: the vehicle's heat	presence, count, speed	works day and night	sensitive to weather
Acoustic	microphone array hears the noise of passing vehicles	presence, count, speed	passive and cheap	less accurate; affected by background noise

Point-to-point

Sensor	How it works	Data	Pros	Cons
Bluetooth / Wi-Fi readers	record anonymised device IDs at two sites and match them	travel time between the sites	cheap; easy to deploy	only a few percent of vehicles are matched; phones now randomise their IDs; privacy
Toll-tag readers	read electronic toll tags (for example E-ZPass) at two sites	travel time	very high match rates where tags are common	only tagged vehicles; needs reader infrastructure
Licence-plate cameras	read plates at two sites and match them	travel time, origin–destination patterns	nearly every vehicle can be matched	privacy and legal limits; misreads; cost

Mobile and aerial

Sensor	How it works	Data	Pros	Cons
GPS probes	phones and fleet vehicles report their positions	speeds and travel times on most roads	wide coverage with no roadside hardware	sample bias and low penetration on minor roads; data licences; privacy

Sensor	How it works	Data	Pros	Cons
Connected vehicles	vehicles broadcast position, speed, acceleration and braking several times a second	detailed vehicle trajectories and driving behaviour	rich, high-frequency data	few equipped vehicles so far; standards and security
Drones	a camera films the road from above, then vehicles are tracked	full trajectories of every vehicle in view	the whole time-space diagram	short flights; regulations and weather; heavy processing
Aerial photographs	a camera on an aircraft photographs the road from straight above	every vehicle in view: density, and speeds from two photos a few seconds apart	a whole stretch at one instant; a classic, simple method	one moment only; flights cost money; weather and daylight
Floating car	a test vehicle drives with the traffic and records its trip	travel time, delay, stops	simple and direct	one vehicle's experience; labour-intensive

Take one tab per group and ask: what slice of the time-space diagram does this sensor see? Which of these could measure a space-mean speed directly? (Drones, and probe vehicles sampled in time.) Which could not see a stopped queue? (Doppler radar.) In the New York region, toll-tag readers have long been used to measure travel times on highways, which makes a good local example.

Seeing from above: aerial photos and drones

From an aircraft



Interstate 64 in Louisville, Kentucky. Every car and truck can be counted. One photo gives the **density** on each lane directly; two photos a few seconds apart give every vehicle's **speed**, and so the space-mean speed. *Photo: NOAA (public domain).*

From a drone



Downtown Athens, filmed by one of ten drones hovering over the city in the pNEUMA experiment (EPFL, 2018). Like an aerial photo, a drone sees every vehicle at once, but it keeps watching: from the video, every trajectory can be recovered. *Video: pNEUMA Vision dataset, Kim, Anagnostopoulos, Barmounakis and Geroliminis (EPFL), CC BY 4.0.*

Compare the two views. An aerial photograph is one instant: a vertical line on the time-space diagram, good for density and, with a second photo, for space-mean speed. Aerial photography was the classic way to measure traffic from above long before drones. The drone keeps filming, so it records the whole time-space diagram: every trajectory. Later in the lecture the same video comes back, with every vehicle tracked by computer vision.

Part 2 · One road, two answers

The puzzle

A loop detector is buried in the road. A drone hovers overhead. Both watch exactly the same cars, and both report the average speed. They will not agree, and neither of them is broken.

i Before you watch: predict

Which sensor will report the **higher** average speed: the loop in the road, or the drone in the sky? Or will they agree? Make a guess before you go on.

The sensor showdown

Two lanes: fast cars spaced out in the top lane, slow cars bunched together in the bottom lane. The orange bar is the loop detector; it flashes each time a car crosses it. The dashed frame is what the drone sees.

Interactive animation: see the web version of these notes. A loop detector reports an average of 85 km/h. A drone watching the same two lanes reports 76 km/h.

What just happened?

- The loop settles near **85 km/h**. The drone reports **76 km/h**. Same cars, same moment, a 9 km/h disagreement.
- **The loop counts fast cars more often**, because fast cars pass it more often.
- **The drone counts slow cars more often**, because slow cars spend longer on the road and bunch up.
- Neither sensor is wrong. They answer **different questions**, and a traffic engineer must know which question the data answered.

Fast sub-stream: 30 m/s, 60 m apart



0.50 cars per second pass a loop, but only 6 fill this stretch: more weight at a point

Slow sub-stream: 15 m/s, 40 m apart



only 0.375 cars per second pass a loop, yet 9 fill this stretch: more weight in a snapshot

Watching...	Like	Average speed
one point , over a period of time	loop detectors, radar guns, cameras	time-mean speed $\bar{v}_t$
a stretch of road , at one instant	drones, aerial photos	space-mean speed $\bar{v}_s$

You drive 30 km to work at 30 km/h, and drive home at 60 km/h. What was your average speed?

Not 45 km/h. The trip took 1 h there and 0.5 h back: 60 km in 1.5 h is **40 km/h**. You spent more *time* driving slowly, so the slow leg counts for more. That is exactly the drone's point of view.

Build your own road

Two lanes, with no lane changing. Set each lane's speed and spacing, and see what each sensor reports.

Interactive figure: see the web version of these notes. With lane A at 30 m/s and 60 m spacing and lane B at 15 m/s and 40 m spacing, a loop detector reports a time-mean speed of 23.6 m/s, while a drone reports a space-mean speed of 21.0 m/s.

Part 3 · Wardrop's theory of the two averages

Wardrop's model (1952)

In 1952 John Glen Wardrop, of Britain's Road Research Laboratory, gave the two averages their precise meaning (*Some theoretical aspects of road traffic research*, Proceedings of the Institution of Civil Engineers).

- The stream is made of M **sub-streams**. In sub-stream i , every vehicle travels at speed v_i , with flow q_i . The total flow is $Q = \sum_i q_i$.
- **In time**, at a fixed point, the share of vehicles with speed v_i is its share of the flow, $f_i = q_i/Q$. The time-mean speed is **flow-weighted**:

$$\bar{v}_t = \sum_i f_i v_i = \frac{\sum_i q_i v_i}{Q}.$$

- **In space**, vehicles of sub-stream i are v_i/q_i metres apart, so their density is $k_i = q_i/v_i$, with $K = \sum_i k_i$. In a snapshot, the share with speed v_i is $f'_i = k_i/K$. The space-mean speed is **density-weighted**:

$$\bar{v}_s = \sum_i f'_i v_i = \frac{\sum_i k_i v_i}{K}.$$

Same traffic, two different speed distributions. Fast vehicles pass a point more often, so they are over-represented in time. Slow vehicles are packed more closely, so they are over-represented in space. The next slide lets students watch the two distributions separate.

Watch it: speeds in time vs speeds in space

A three-lane road where every driver keeps their own speed. The loop records the speed of each car that passes; the drone photographs the road once a second. Watch the two histograms separate.

Interactive animation: see the web version of these notes. The loop's histogram of speeds leans towards fast cars; the drone's leans towards slow cars. The harmonic mean of the loop's speeds matches the drone's plain average.

Three things to point out. The drone histogram sits to the left of the loop histogram. The harmonic mean of the loop's speeds agrees with the drone's plain average, which is the "which average for which data" table in action. And with the spread set to zero, all three numbers agree.

And $q = kv$ falls out

Since $k_i v_i = q_i$, the numerator of the space-mean speed is just Q :

$$\begin{aligned}\bar{v}_s &= \frac{\sum_i q_i}{K} = \frac{Q}{K} \\ \Rightarrow Q &= K \bar{v}_s.\end{aligned}$$

The fundamental relation $q = kv$ holds with the **space**-mean speed, and only with it.

Which average for which data?

Wardrop's definitions tell us how to compute each mean from the data we actually have:

Data collected...	Time-mean speed $\bar{v}_t$	Space-mean speed $\bar{v}_s$
at a point (N vehicles passing, speeds v_j)	arithmetic mean: $\frac{1}{N} \sum_j v_j$	harmonic mean: $\frac{N}{\sum_j 1/v_j}$
in space (N vehicles in a snapshot)	$\frac{\sum_j v_j^2}{\sum_j v_j}$	arithmetic mean: $\frac{1}{N} \sum_j v_j$

! A common misconception

“The arithmetic mean gives the time-mean speed, and the harmonic mean gives the space-mean speed.” **Not in general.** Which average you need depends on **where the data were collected**. The plain average of speeds from a drone photo *is* the space-mean speed. The harmonic mean gives the space-mean speed only for data collected at a point.

The harmonic mean has a natural reading. Multiply its numerator and denominator by a section length L : $\bar{v}_s = L/(\frac{1}{N} \sum_j L/v_j)$, the section length divided by the **average travel time**. That is why the space-mean speed is the one to use for travel time.

The proof, part 1: rewrite the time-mean speed

We will prove $\bar{v}_t = \bar{v}_s + \sigma_s^2/\bar{v}_s$.

Step 1. Write the time-mean speed in space terms. Substitute $q_i = k_i v_i$:

$$\begin{aligned}\bar{v}_t &= \frac{\sum_i q_i v_i}{\sum_i q_i} \\ &= \frac{\sum_i k_i v_i^2}{\sum_i k_i v_i}.\end{aligned}$$

Step 2. Divide the top and bottom by K . Both become averages over the space distribution $f'_i = k_i/K$:

$$\begin{aligned}\bar{v}_t &= \frac{\sum_i f'_i v_i^2}{\sum_i f'_i v_i} \\ &= \frac{\sum_i f'_i v_i^2}{\bar{v}_s}.\end{aligned}$$

The proof, part 2: expand around the mean

Step 3. Expand around the mean. Write $v_i = \bar{v}_s + (v_i - \bar{v}_s)$ and square:

$$\begin{aligned}\sum_i f'_i v_i^2 &= \bar{v}_s^2 + 2 \bar{v}_s \underbrace{\sum_i f'_i (v_i - \bar{v}_s)}_{=0} \\ &\quad + \underbrace{\sum_i f'_i (v_i - \bar{v}_s)^2}_{=\sigma_s^2}.\end{aligned}$$

The middle term vanishes because the deviations from a mean average to zero. The last term is the **variance of speeds in space**, σ_s^2 .

Step 4. Put it together.

$$\begin{aligned}\bar{v}_t &= \frac{\bar{v}_s^2 + \sigma_s^2}{\bar{v}_s} \\ &= \bar{v}_s + \frac{\sigma_s^2}{\bar{v}_s}\end{aligned}$$

What it tells us:

- The time-mean speed is **never smaller** than the space-mean speed, because a variance cannot be negative.
- They are equal **only when every vehicle has the same speed** ($\sigma_s = 0$).
- The gap grows with the **spread** of speeds, and it is largest in slow, mixed traffic, where $\bar{v}_s$ is small.
- The variance must be the one measured **in space**. Using the variance of speeds measured at a point gives the wrong answer.

A worked example

i Check it with numbers

Lane A: 24 m/s, 48 m apart. Lane B: 12 m/s, 24 m apart.

- Flows: $q_A = 24/48 = 0.5$ veh/s and $q_B = 12/24 = 0.5$ veh/s, so $\bar{v}_t = (0.5 \times 24 + 0.5 \times 12)/1 = 18$ m/s.
- Densities: $k_A = 1/48$, $k_B = 1/24$ veh/m, so $f'_A = 1/3$, $f'_B = 2/3$ and $\bar{v}_s = 24/3 + 2 \times 12/3 = 16$ m/s.
- Space variance: $\sigma_s^2 = \frac{1}{3}(24 - 16)^2 + \frac{2}{3}(12 - 16)^2 = 32$, so $\bar{v}_s + \sigma_s^2/\bar{v}_s = 16 + 2 = 18$ m/s.

Enter these numbers in *Build your own road* above to check them live.

Edie's definitions: try any region

Take **any region** A of the time-space diagram, of area $|A|$. Add up the total distance $d(A)$ and the total time $t(A)$ that vehicles spend inside it:

$$\begin{aligned}q(A) &= \frac{d(A)}{|A|}, \\ k(A) &= \frac{t(A)}{|A|}, \\ v(A) &= \frac{d(A)}{t(A)}.\end{aligned}$$

Move the region around the red-light diagram, and change its shape: $q = kv$ holds every time.

Interactive figure: see the web version of these notes. Edie's definitions give flow, density and speed for any region of the time-space diagram, and $q = kv$ holds exactly for every region.

Try the three shapes. The thin horizontal region is what a loop detector sees: its q is the count per hour, and its v is the harmonic mean of passing speeds. The thin vertical region is a drone snapshot: its k is the number of cars per kilometre, and its v is their plain average. The square is neither, and Edie's definitions still work. Put the square in the queue during the red: k is high, q and v are near zero.

Part 4 · Sensors in the real world

Occupancy is a space-mean measurement

Where the estimate comes from. Vehicle j , of length ℓ_j and speed $\dot{x}_j$, covers a loop of length d for $\tau_j = (d + \ell_j)/\dot{x}_j$ seconds. Over a period T with N vehicles:

$$\begin{aligned} o &= \frac{1}{T} \sum_{j=1}^N \tau_j = \frac{1}{T} \sum_{j=1}^N \frac{d + \ell_j}{\dot{x}_j} \\ &\approx (d + \ell) \frac{1}{T} \sum_{j=1}^N \frac{1}{\dot{x}_j} \quad (\text{all } \ell_j \approx \ell) \\ &= (d + \ell) \underbrace{\frac{N}{T}}_q \underbrace{\frac{1}{N} \sum_{j=1}^N \frac{1}{\dot{x}_j}}_{1/\bar{v}_s} \\ &= (d + \ell) \frac{q}{\bar{v}_s} = (d + \ell) k. \end{aligned}$$

So $k \approx o/(d + \ell)$. Notice the **harmonic mean** of the passing speeds: occupancy is secretly a space-mean measurement. The weak step is $\ell_j \approx \ell$, the assumption that every vehicle has the same length.

Think about it: occupancy is a fraction with no units, while k is in vehicles per kilometre. Where do the units come from in $k = o/(d + \ell)$?

Live: a loop detector, and trucks that fool it

A single lane at a steady 72 km/h passes a 1.8 m loop. The loop flashes while a vehicle is over it. Every 30 seconds it reports what a real detector would: a count, an occupancy, and a speed estimated by assuming every vehicle is a 5 m car.

Interactive animation: see the web version of these notes. With 20% trucks, a single loop that assumes 5 m cars reports a speed about 28% too low.

Set trucks to 0%: the estimate matches 72 km/h, apart from random noise within each 30-second period. At 20% the estimate drops to about 55 km/h, although nobody slowed down. That is why real systems use dual loops or radar to measure speed directly, and why a speed drop at a single loop is not always congestion.

Your phone is a sensor too


Navigation apps collect GPS “pings” from phones at regular intervals. A car that spends twice as long on a stretch sends twice as many pings, so averaging the pings gives a **space-mean speed**, just like the drone. The catch: only cars running the app are counted.

Interactive figure: see the web version of these notes. With 5% of cars sending pings, the app’s space-mean speed is typically off by about 10 km/h. More users shrink the scatter, but if app users drive faster than everyone else, a bias remains however many there are.

Two lessons. Low penetration gives **noise**: the scatter shrinks as the share of app users grows. A biased sample gives **bias**: set “drive faster” to 10% and push penetration to 50%, and the cloud narrows but stays off-centre. More data cures noise, never bias.


Why it matters: a fifth of your travel time

On real US-101 traffic (NGSIM, 8:00 a.m., the same four minutes you will load in Lab 1), a virtual loop detector reports a time-mean speed of **44.6 km/h**. Edie’s space-mean speed for the same traffic is **35.0 km/h**.



Use the loop's number to predict travel time, and you **underestimate it by about a fifth**. Every travel-time sign and every navigation app faces the same choice of averages.


AI at work: from raw data to clean data



The Athens drone video with every vehicle detected and tracked by computer vision (pNEUMA Vision, EPFL, CC BY 4.0).

- **Collecting:** deep-learning detectors and trackers turn video into trajectories almost automatically.
- **Cleaning:** anomaly detection flags broken detectors: stuck at zero, frozen, or physically impossible.
- **Filling gaps:** imputation, from regression to neural networks, estimates missing values.
- **The catch:** each step is itself a model, and models can be wrong.

the middle step is a model: it can be wrong



A tracker that swaps two cars’ identities, or an imputation that fills a gap with a smooth “normal day”, changes the traffic you think you see. In PEMS08, a public freeway dataset widely used to benchmark AI forecasting, some detectors report exactly identical time series: a sign that missing data were filled in before release. Always ask how your data were made. NGSIM’s trajectories, by contrast, took years of work in the mid-2000s, with tracking software followed by manual checking. You will do both jobs yourself. In Lecture 3 you will see a deep-learning detector and tracker turn drone video into trajectories; in Lab 3 you clean such tracks and check them against physics, and Lab 4 hunts for broken detectors and fills their gaps.

Part 5 · Beyond counting cars

LiDAR at the intersection

A **LiDAR** sensor sweeps laser pulses around itself and times each echo. Many times a second it produces a **point cloud**: a 3-D picture of everything in view, by day or night.

- **From points to trajectories.** Remove the background (road, buildings, trees), group the remaining points into objects, label each object a vehicle or a pedestrian, then follow it from frame to frame (Wu et al., *TRR*, 2018; Zhao et al., *TR-C*, 2019). The result is the trajectory of **every road user**: position, speed and direction.
- **From the lab to the street.** In 2017 Hao Xu’s group at the University of Nevada, Reno mounted a LiDAR beside Virginia Street in Reno, one of the first roadside uses of the sensor. More followed along the corridor and in Henderson, Nevada.
- **Privacy by design.** The data are only points: no colour, no licence plates, no faces (UNR, 2020). NCHRP’s guide to signal-based counting notes that, unlike cameras, LiDAR raises no data-privacy concerns (NCHRP Web-Only Document 436, 2025).

This is Lecture 1’s trajectories again, now collected from a pole at the corner instead of a drone or a tall building. Ask the class which slice of the time-space diagram a roadside LiDAR sees: every trajectory, but only within its range, typically a single intersection.

Watching people, not just cars

A loop detector cannot see a pedestrian, and a push button counts requests, not people. Yet pedestrians and cyclists are the road users most at risk.

- **How people really cross.** LiDAR at intersections in Texas and Utah tracked pedestrians and found that their crossing paths often differ from what the design assumed. The same project built a “dynamic flashing yellow arrow” that reacts to pedestrians the LiDAR sees about to cross (National Institute for Transportation and Communities, 2023).
- **Seeing through a crowd.** A Caltrans-funded project at UC Riverside fused camera and LiDAR data to detect pedestrians and cyclists even when they are partly hidden, with over 90% recall, and to warn connected vehicles of them (Caltrans, 2025).

Safety without waiting for crashes

Crashes are rare. Waiting for enough of them to judge an intersection can take years, and each one costs lives. **Conflicts**, or near misses, are far more common, and trajectories reveal them.

- **Time-to-collision (TTC):** how long until two road users collide if both keep their speed and path (Hayward, 1972). For a follower closing on its leader with bumper-to-bumper gap g :

$$\text{TTC} = \frac{g}{v_{\text{follower}} - v_{\text{leader}}},$$
$$v_{\text{follower}} > v_{\text{leader}}.$$

- **Post-encroachment time (PET):** the time between one road user leaving a conflict point and the next arriving there (Allen, Shin and Cooper, 1978).
- **Measured three ways:** from video (Ismail et al., *TRR*, 2009), from roadside LiDAR, including vehicle–pedestrian near-crashes (Wu et al., *Accident Analysis & Prevention*, 2018), and from simulation, with FHWA’s Surrogate Safety Assessment Model counting a conflict when $\text{TTC} < 1.5$ s or $\text{PET} < 5$ s by default (Gettman et al., 2008).

Connect this to the time-space diagram. For two cars in one lane, the gap is the vertical distance between their trajectories, and the closing speed is the difference in their slopes. TTC is how long until the two lines would meet. A small TTC means two trajectories almost touching. This is why complete trajectories, from drones or LiDAR, are so valuable for safety.

Road inspection: sensors on the move

Sensors do not only watch traffic. Mounted on a van driving at traffic speed, they inspect the road itself.

- **Mobile LiDAR mapping.** Scanners, cameras and precise positioning (GNSS and inertial sensors) map lane markings, signs, poles, manholes and bridge clearances in one pass (Guan et al., 2016). NCHRP Report 748 gives state DOTs guidelines for using it (Olsen et al., 2013).
- **Pavement condition.** 3-D laser systems scan a strip 4 m wide, with about 4,000 points across the lane, vertical accuracy of about ± 0.5 mm and a profile every 5 mm at 100 km/h, day or night. From these scans agencies measure cracking, rutting and faulting (FHWA, 2020).

One sensor is never enough: fusion

Sensor	Strong at	Weak at
Camera	detail; telling a bus from a truck from a cyclist	darkness, glare, rain, snow, fog
Radar	any weather, day or night; long range	coarse detail; poor at telling objects apart
LiDAR	precise 3-D positions; unaffected by light	cost, occlusion, heavy data; rain, snow and fog add noise
Loop detector	reliable counts and occupancy at one point	one point only; cannot see pedestrians
Phone and vehicle GPS	whole trips, citywide	only some vehicles, and not a random sample

Sources: Yeong et al. (*Sensors*, 2021), Sun et al. (*Sensors*, 2022), NCHRP Web-Only Document 436 (2025).

Data fusion combines sensors so each covers the others’ blind spots: loops for counts plus probes for travel times, or cameras plus LiDAR for pedestrians (El Faouzi, Leung and Kurian, *Information Fusion*, 2011). Like the AI cleaning on the previous slide, fusion is itself a model, so it inherits the same question: how were these data made?

Further reading: sensing beyond counting cars

- Allen, B. L., Shin, B. T. and Cooper, P. J. (1978). Analysis of traffic conflicts and collisions. *Transportation Research Record* 667, 67–74. <https://onlinepubs.trb.org/Onlinepubs/trr/1978/667/667-009.pdf>
- Caltrans (2025). Wei, C., Zhang, Z. and Wu, G. *Vulnerable Road User (VRU): Accurate and Reliable Detection for Roadside-Assisted Cooperative Driving*. Report CA-27-4067. <https://dot.ca.gov/-/media/dot-media/programs/research-innovation-system-information/documents/final-reports/ca-27-4067-a11y-v2.pdf>
- Chang, G., Gilliland, A., Rada, G. R., Serigos, P. A., Simpson, A. L. and Kouchaki, S. (2020). *Successful Practices for Quality Management of Pavement Surface Condition Data Collection and Analysis*. FHWA-RC-20-0007. <https://www.fhwa.dot.gov/pavement/management/pubs/fhwarc20007.pdf>
- El Faouzi, N.-E., Leung, H. and Kurian, A. (2011). Data fusion in intelligent transportation systems: Progress and challenges – A survey. *Information Fusion* 12(1), 4–10. <https://doi.org/10.1016/j.inffus.2010.06.001>
- Gettman, D., Pu, L., Sayed, T. and Shelby, S. (2008). *Surrogate Safety Assessment Model and Validation: Final Report*. FHWA-HRT-08-051. <https://www.fhwa.dot.gov/publications/research/safety/08051/>
- Guan, H., Li, J., Cao, S. and Yu, Y. (2016). Use of mobile LiDAR in road information inventory: A review. *International Journal of Image and Data Fusion* 7(3), 219–242. <https://doi.org/10.1080/19479832.2016.1188860>
- Hayward, J. C. (1972). Near-miss determination through use of a scale of danger. *Highway Research Record* 384, 24–34. <https://onlinepubs.trb.org/Onlinepubs/hrr/1972/384/384-004.pdf>
- Ismail, K., Sayed, T., Saunier, N. and Lim, C. (2009). Automated analysis of pedestrian–vehicle conflicts using video data. *Transportation Research Record* 2140, 44–54. <https://doi.org/10.3141/2140-05>
- Mahmud, S. M. S., Ferreira, L., Hoque, M. S. and Tavassoli, A. (2017). Application of proximal surrogate indicators for safety evaluation: A review of recent developments and research needs. *IATSS Research* 41(4), 153–163. <https://doi.org/10.1016/j.iatssr.2017.02.001>
- National Institute for Transportation and Communities (2023). *Pedestrian Behavior Study to Advance Pedestrian Safety in Smart Transportation Systems Using Innovative LiDAR Sensors*. NITC-RR-1393. <https://trec.pdx.edu/news/new-lidar-system-pinpoints-pedestrian-behavior-improve-efficiency-and-safety-intersections>
- Olsen, M. J., Roe, G. V., Glennie, C., Persi, F., Reedy, M., Hurwitz, D., Williams, K., Tuss, H., Squellati, A. and Knodler, M. (2013). *Guidelines for the Use of Mobile LIDAR in Transportation Applications*. NCHRP Report 748. https://onlinepubs.trb.org/onlinepubs/nchrp/nchrp_rpt_748.pdf
- Sun, P., Sun, C., Wang, R. and Zhao, X. (2022). Object detection based on roadside LiDAR for cooperative driving automation: A review. *Sensors* 22(23), 9316. <https://doi.org/10.3390/s22239316>
- Tsapakis, I., et al. (2025). *Leveraging Existing Traffic Signal Assets to Obtain Quality Traffic Counts: A Guide*. NCHRP Web-Only Document 436. <https://doi.org/10.17226/29214>
- University of Nevada, Reno (2020). New roadside LiDAR sensors help build a safer transportation infrastructure. <https://www.unr.edu/nevada-today/news/2020/hao-xu-roadside-lidar>
- Wu, J., Xu, H., Sun, Y., Zheng, J. and Yue, R. (2018). Automatic background filtering method for roadside LiDAR data. *Transportation Research Record* 2672(45), 106–114. <https://doi.org/10.1177/0361198118775841>
- Wu, J., Xu, H., Zheng, Y. and Tian, Z. (2018). A novel method of vehicle-pedestrian near-crash identification with roadside LiDAR data. *Accident Analysis & Prevention* 121, 238–249. <https://doi.org/10.1016/j.aap.2018.09.001>
- Yeong, D. J., Velasco-Hernandez, G., Barry, J. and Walsh, J. (2021). Sensor and sensor fusion technology in autonomous vehicles: A review. *Sensors* 21(6), 2140. <https://doi.org/10.3390/s21062140>
- Zhao, J., Xu, H., Liu, H., Wu, J., Zheng, Y. and Wu, D. (2019). Detection and tracking of pedestrians and vehicles using roadside LiDAR sensors. *Transportation Research Part C* 100, 68–87. <https://doi.org/10.1016/j.trc.2019.01.007>

Part 6 · Wrap-up

Three challenges

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

1. **Sensor truce.** In *Build your own road*, find settings where the loop and the drone agree exactly. What must be true of the traffic?
2. **Maximum disagreement.** Make the loop report at least **twice** the drone's speed. Describe the traffic you built. Where might you see it in real life?
3. **Truck trouble.** What share of trucks makes a single loop's speed estimate 25% too low?

Summary

- “Average speed” has two answers: **time-mean** (point sensors) and **space-mean** (snapshots, and phone pings).
- Wardrop (1952): the time-mean speed is **flow-weighted**, the space-mean speed is **density-weighted**, and $q = k \bar{v}_s$.
- We proved $\bar{v}_t = \bar{v}_s + \sigma_s^2 / \bar{v}_s$, so the time-mean speed is never the smaller.
- Which average to compute depends on **where the data were collected**, not on which speed you want.
- Edie's definitions work for any region, and $q = kv$ always holds with them.
- Picking the wrong average can cost a fifth of your travel-time estimate, and a single loop can be fooled by trucks.
- Machine learning now collects, cleans and fills traffic data, but each of those steps is a model that can be wrong.
- Sensors sit in the pavement, beside the road, in vehicles or in the air, and see a point, two points, an area or a trajectory.
- LiDAR and multi-sensor systems see pedestrians, near misses and the road itself. No sensor sees everything, so agencies fuse them.
- Every sensor sees traffic through its own window. **Know which window your data came from.**
- **Next week:** the fundamental diagram, the most important curve in traffic, and the puzzles it cannot explain.

Notation

Symbol	Meaning	Units
$\bar{v}_t$	time-mean speed: the average over vehicles passing a point	m/s or km/h
$\bar{v}_s$	space-mean speed: the average over vehicles on a stretch of road	m/s or km/h
σ_s^2	variance of speeds measured in space	(m/s) ²
M	number of sub-streams in Wardrop's model	—
q_i, k_i, v_i	flow, density and speed of sub-stream i	veh/h, veh/km, km/h
Q, K	total flow and total density	veh/h, veh/km
f_i, f'_i	share of sub-stream i in time (q_i/Q) and in space (k_i/K)	—
$A, A $	a region of the time-space diagram, and its area	—, m · s

Symbol	Meaning	Units
$d(A), t(A)$	total distance and total time travelled by vehicles inside A	m, s
N	number of vehicles observed	veh
o	occupancy: fraction of time the loop is covered	—
d	loop length (not the same as Edie's $d(A)$)	m
ℓ, ℓ_j	assumed vehicle length; length of vehicle j	m
τ_j	time vehicle j covers the loop	s
T	length of the counting period	s
g	bumper-to-bumper gap between a follower and its leader	m
TTC	time-to-collision: time until two road users collide if neither changes speed or path	s
PET	post-encroachment time: time between one road user leaving a conflict point and the next arriving	s