Line of Travel



Learns the order Amazon drivers actually deliver in, then removes only the detours.

−5.9%
drive time vs the driver's own order

0.01
zone re-entries per route (shortest path: 15.7)

3,052
real held-out routes in 5 US metros

SANCHAY SINGH · SOLO ENTRY · LIVE DEMO: HTTPS://SANCHAY117.GITHUB.IO/LINE-OF-TRAVEL/

The shortest route is 11.3% shorter. Drivers won't drive it.

DRIVER'S ORDER



DRIVE TIME	ZONE RE-ENTRIES	AMAZON SCORE
3h 41m	1	0.000

SHORTEST PATH



DRIVE TIME	ZONE RE-ENTRIES	AMAZON SCORE
3h 10m −14.0%	23	0.025

LINE OF TRAVEL



DRIVE TIME	ZONE RE-ENTRIES	AMAZON SCORE
3h 21m −9.1%	0	0.002

A pure optimizer ignores the zones drivers work in. On the held-out routes it re-enters 15.7 zones per route, where drivers re-enter 1.4. A plan the driver will not follow saves nothing, which is why Amazon's 2021 challenge scored routes on how closely they match what drivers really do.

Drivers work zone by zone, in an order they repeat.

ZONES KEPT WHOLE

94%

of zones the driver finishes in one visit (about 19 zones per route)

NEXT ZONE PREDICTED

68%

of the driver's next zones our model predicts, from about 21 choices

NEAREST-ZONE GUESS

35%

if you simply pick the nearest zone



Where the name comes from. Mail carriers sort each day's mail into a case of pigeonholes in the order they walk. The US Postal Service records that order for every carrier route and calls it the Line of Travel. Amazon drivers have one too. Nobody wrote it down, so we learned it.

Built only on Amazon's own routes, start to finish.

01 · DATA

Amazon Last Mile Routing Research Challenge 2021

9,164 routes, 1,337,758 stops, 17 stations in 5 metros. Public S3, CC BY-NC 4.0.

02 · INGEST

Parquet tables and float32 travel matrices

Stops, zones, packages, delivery windows, service times, the driver's actual order.

03 · LEARN

Next-zone ranker (LightGBM)

Trained on the 6,112 training routes. Gives every zone-to-zone move a habit cost.

04 · SOLVE

Driver-aware sequencing engine

Plain JavaScript. The same file runs the benchmark in Node and the demo in your browser.

05 · SCORE

Held-out benchmark

3,052 evaluation routes, Amazon's official route score, bootstrap intervals.

6,112 / 3,052

training / held-out routes, split by Amazon

6.6%

of held-out stops carry a delivery time window

n × n

travel seconds per route, Amazon's road-network estimates

0

personal data used: no driver IDs, no customer names

METHOD

Price every move the way a driver would. Then search.

EDGE COST USED BY THE SEARCH

$$\text{cost}(i \to j) = \text{travel}(i, j) + [j \text{ in a new zone}] \times (\text{crossing} + w \times -\log p(\text{zone } j \mid \text{zone } i))$$

travel(i, j) — Seconds on the road, straight from Amazon's travel-time matrix for that route.

crossing — A flat charge for leaving a zone, so the van finishes a zone before it moves on.

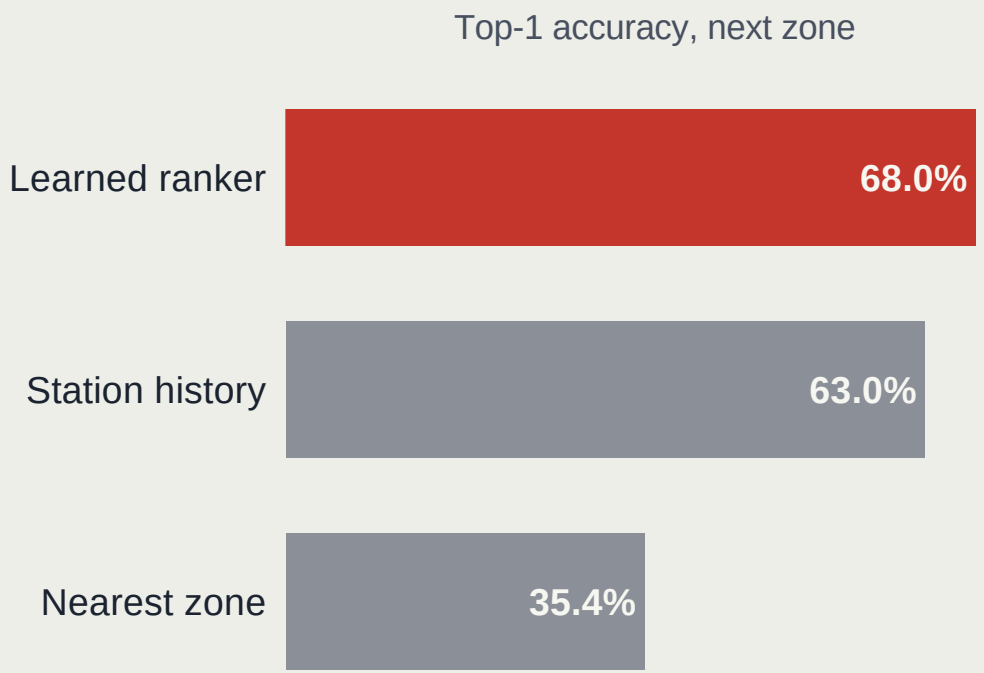
−log p — The habit cost: cheap for moves drivers usually make, expensive for moves they rarely make.

THE SEARCH, 0.3 S PER ROUTE IN THE BENCHMARK

1. Start from nearest neighbour under that cost.
2. Improve with asymmetric 2-opt, or-opt, and moves that lift a whole zone to a better place.
3. Kick with double bridges, half of them cut at zone boundaries.
4. Accept on cost plus 20 × seconds late, from a full schedule simulation.
5. Repair windows: move whole zones first, single stops only if still late.

TRUST DIAL OFF: TRAVEL TIME ONLY (SHORTEST PATH) · ZONES ONLY: W 0, CROSSING 600 S · MEDIUM: W 30, CROSSING 300 S · HIGH: W 60, CROSSING 600 S · MAXIMUM: W 300, CROSSING 1200 S

The model knows where drivers go next.



66,621 HELD-OUT ZONE-TO-ZONE MOVES

About 21 zones to choose from at each step. Chance would be under 5%.

HOW IT WORKS

- A LightGBM lambdarank model ranks every unvisited zone as the next one.
- Features: distance between zone centres and its rank among the candidates, steps between Amazon's zone codes (cluster, block, row, column), and how often drivers at this station made the same move on other days.
- Station history for a training route counts other days only, so the model never sees the answer.
- Scores become a habit cost of −log p for each move, which the search adds to travel time.

5.9% less driving, and closer to the driver than any optimizer.

METHOD	DRIVE TIME VS DRIVER	95% INTERVAL	AMAZON SCORE (0 = DRIVER)	ZONE RE-ENTRIES PER ROUTE	LATE STOPS
Driver's own order	baseline		0	1.38	469
Nearest neighbour	+2.2%		0.105	20.98	2,027
Shortest path (our engine, travel time only)	−11.3%	11.2 to 11.5%	0.080	15.72	1
Line of Travel, zones only	−8.0%		0.050	0.00	2
Line of Travel, default	**−5.9%**	**5.8 to 6.1%**	**0.032**	**0.01**	**1**

DRIVER-HOURS

592 h

of driving saved across the 3,052 routes, about 11.6 minutes per route

MEDIAN ROUTE

5.5%

saved on the median route: half of all routes save more than this

ROUTES THAT LOSE

2.4%

of routes drive longer than the driver did; 0.2% by more than 4%

Trust is a dial. Each setting is a different deal with the driver.



Shortest path

−11.3% drive time, but 15.7 zone re-entries per route.

Zones only

−8.0%. Keeps zones whole, ignores the driver's zone order.

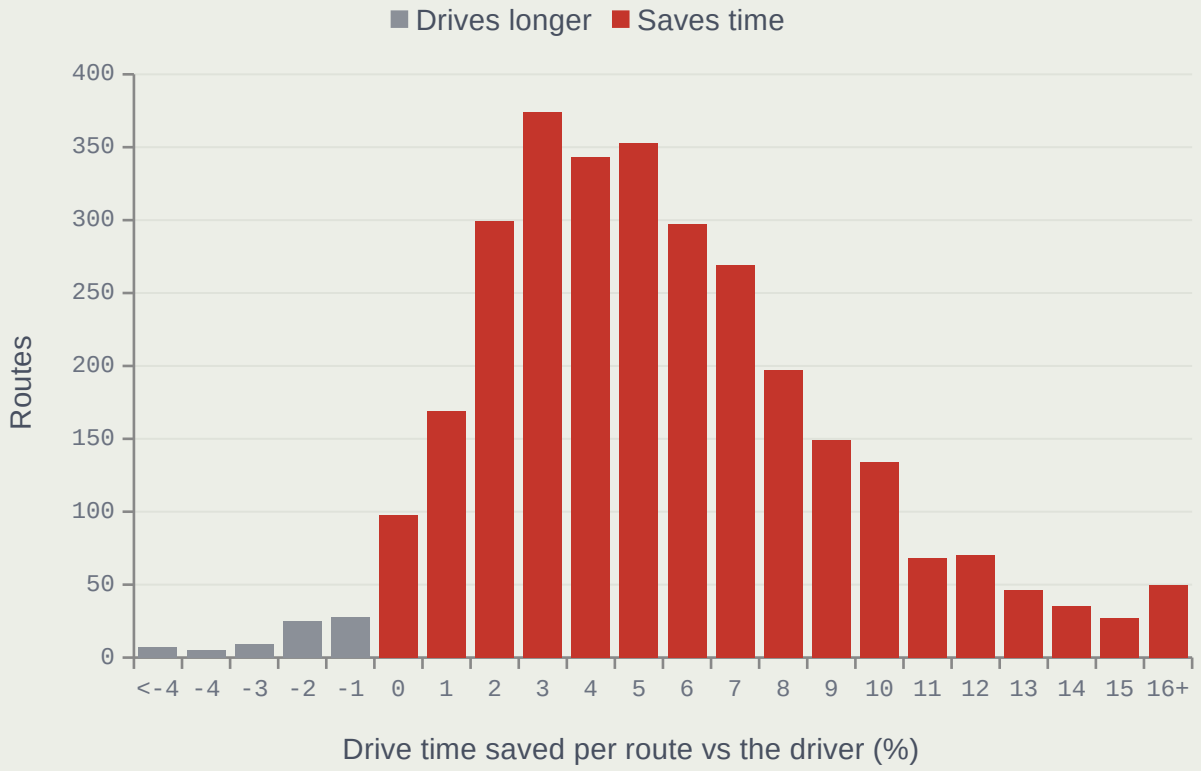
Default

−5.9%. Keeps zones whole and follows the usual order. Score 0.032.

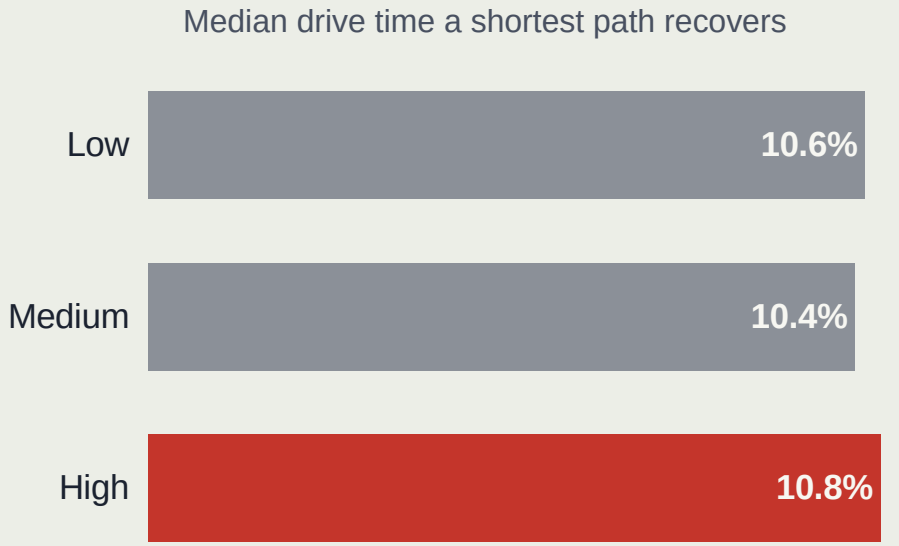
High and max trust

−5.2% and −4.3%. For stations that want plans closest to habit.

It holds route by route, and the waste is on well-rated routes too.



3,052 held-out routes at the default setting. Median 5.5%. 2.4% of routes drive longer than the driver did.



Even routes Amazon rated High have about 11% to recover. Median drive time a shortest path recovers, by Amazon's own route-quality label, 6,112 training routes. The labels do not predict the waste, so the saving is not confined to a few badly run routes.

A late van keeps its saving. Rush orders keep every deadline.

SCENARIO	WHAT CHANGES	DRIVE TIME, LINE OF TRAVEL	DRIVE TIME, SHORTEST PATH	LATE STOPS: DRIVER	LATE STOPS: LINE OF TRAVEL	RE-ENTRIES, LINE OF TRAVEL
Normal day	The day as planned.	**−6.2%**	−11.6%	61	**0**	0.01
Van leaves 45 min late	Van leaves 45 minutes late, so every window closes 45 minutes sooner.	**−5.7%**	−11.2%	227	**0**	0.01
8 same-day orders added	8 same-day orders near existing stops, each due within 4 hours. Driver inserts each where it adds least time.	**+5.3%**	−4.9%	1,809	**0**	1.75

LATE START

227 → 0

late stops after a 45-minute late start; drive time still −5.7%

RUSH ORDERS

1,809 → 0

missed deadlines with 8 rush orders, for +5.3% drive time

ZONES KEPT

1.75 vs 20.5

zone re-entries per route with rush orders: ours vs shortest path

Every number is reproducible from one script.

Held-out test. The habit model learns from Amazon's 6,112 training routes. Every number here is scored on the 3,052 evaluation routes Amazon held back.

Official metric. Amazon score uses a port of the challenge scorer that matches the original exactly on 40 randomised test routes.

Solver quality. On 200 routes our engine's shortest path (the code the browser runs) is 2.26% shorter than Google OR-Tools given 5 s per route, with a 0.3 s budget.

Ablation. Our first engine saved 3.3% and drove longer than the driver on 18% of routes. Moving whole zones in the search and counting lateness when accepting a tour took that to 5.9% and 2.4%.

Uncertainty. Drive time saved at the default setting: 5.9% (95% bootstrap interval 5.8% to 6.1%). Shortest path: 11.3% (11.2% to 11.5%).

Delivery windows. 28,354 stops have a time window. Late stops: driver's order 469, Line of Travel 1.

No leakage. Station-history features for a training route only count other days; evaluation routes only see training history.

Honest limit. Travel times are Amazon's planned estimates, not GPS traces, so savings are measured in Amazon's own matrix. Drivers may know shortcuts the matrix does not.

Pick a real route. Decide how much to trust the driver.

CITY

Austin

ROUTE

DAU1 · Sep 11, 2018 · 183 stops

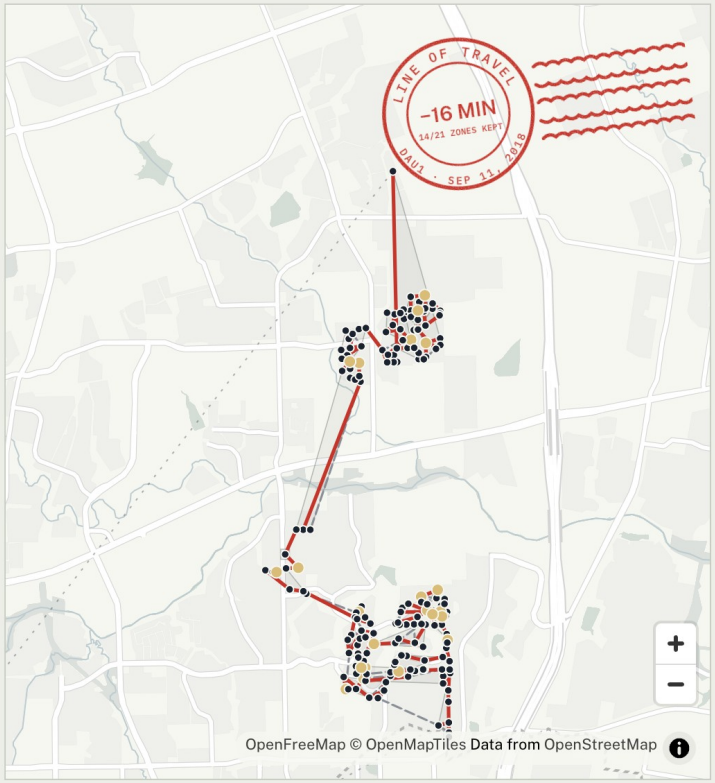
TRUST THE DRIVER MEDIUM

Shortest path — Driver's habits

WHAT IF

- (●) Normal day
- ( ) Van leaves 45 min late
- ( ) 8 same-day orders added

RE-SEQUENCE



SAME STOPS, SAME MATRIX

	DRIVER	LINE OF TRAVEL	SHORTEST PATH
Drive time	3h 12m	2h 56m	**2h 50m**
Day length	7h 01m	6h 45m	**6h 39m**
Late stops	**0**	**0**	**0**
Zone re-entries	**0**	**0**	19
Amazon score	**0.000**	0.013	0.026
Solve time	n/a	578 ms	490 ms

Amazon score is the challenge's own route-similarity metric: 0 means identical to the driver.

TRY IT

- 36 real held-out routes from 5 cities.
- Trust dial from shortest path to the driver's habits.
- What-if: late start, or 8 same-day orders.
- Drive time, day length, late stops, zone re-entries and Amazon score, side by side.
- Solved in your browser by the benchmark's own engine: 0.4 to 0.7 s a route on an M4 Pro laptop. No backend.

OPEN

https://sanchay117.github.io/line-of-travel/

Minutes per van, every day, with nothing new to install.

PER VAN

11.6 min

less driving per route per day at the default setting, on held-out routes

PER STATION

19.4 h

of driving saved per day at a 100-van station

PER YEAR

$102k

a year per station at the BLS 2025 median wage for delivery drivers ($21.13/h), 250 days

HOW IT FITS INTO A STATION

1 Plan

The station's planner assigns stops to vans, as today.

2 Re-sequence

Line of Travel orders each van's stops in well under a second, in a browser.

3 Dial

Each station sets its trust level. New drivers can start nearer the shortest path.

4 Learn

Each day's completed routes add to the station history the model reads.

Fewer late deliveries too: 469 late stops in the drivers' own orders, 1 with Line of Travel, across 28,354 stops with a window. The US has about 1.5 million delivery truck driver jobs (BLS 2025).

It keeps the driver's knowledge. Here is what it cannot know.

WHAT IT PROTECTS

Driver judgment. It starts from how drivers already work instead of overriding it, and a dial decides how far it may depart.

Deliveries. Lateness is part of the objective, so saving minutes never comes from missing a window.

Privacy. Only stops, zones and travel times. No driver identity, no customer names, no GPS traces.

Explainability. Every choice is road time plus two named charges a dispatcher can read and change.

LIMITS

Planned, not driven. Savings are measured in Amazon's travel-time estimates, not GPS. Drivers may know shortcuts the matrix does not.

Habits are not always good. Some habits are detours. That is why the dial exists, and why the default does not trust habit fully.

Five US metros, 2018. Other countries, vehicles or seasons need their own history before the habit model is trusted.

Research licence. The dataset is CC BY-NC 4.0, so a commercial rollout would retrain on the operator's own routes.

Keep the order drivers already know.
Remove the detours.



LIVE DEMO

https://sanchay117.github.io/line-of-travel/

CODE AND REPRODUCTION

https://github.com/Sanchay117/line-of-travel

DATA

Amazon Last Mile Routing Research Challenge 2021, CC BY-NC 4.0

SANCHAY SINGH · SOLO ENTRY · DATANOVA 2026