

Reliable Routing from Sparse Observations

Anil Kumar¹, Suresh Rao², Mahesh Reddy³

Institute of Aeronautical Engineering

Abstract

Sparse observations can make a reliable route difficult to certify even when a feasible path exists. We study routing with unknown link-disruption probabilities and bounded spatial variation. A multiscale estimator pools observations, corrects for between-link bias, and produces simultaneous upper confidence bounds. An exact constrained search then minimizes distance subject to a route-failure limit. The coverage statement permits both pooling radii and routes to be selected after observing the data. We evaluate the method on a directed street network using three synthetic hazard fields, three observation patterns, 24 trip pairs, and four data budgets. Comparisons include direct confidence bounds, a single-observation-source transfer envelope, pooled point estimates, a bias ablation, and a true-model oracle. Spatial pooling substantially increased acceptance in a smooth field, while concentrated hazards reduced its advantage through larger bias allowances. The confidence-based routes produced no observed risk violations; point estimates and omitted bias corrections sometimes exceeded the prescribed risk. These results identify when spatial information helps reliable routing and when incomplete sensing or sharp hazard variation remains limiting.

Introduction

Reliability depends on what is known about every part of a route. An operator may have repeated observations on major corridors and little information about alternate links. A route through an unobserved area can be short and physically usable while remaining difficult to certify from the available data. The planning problem therefore couples a network decision with a statistical question: which observations support the probabilities used along the selected path?

Spatial side information connects learning and planning across related states [1–3]. Bounds on differences between transition probabilities allow observations from one state to constrain estimates at another. Indirect-sampling methods incorporate this information into model learning and exploration [4]. For routing, this requires a quantitative relationship between link similarity and disruption probability. We use that relationship to pool observations while accounting for the mismatch between source and target links.

Reliable routing combines travel-time optimization with probability constraints and

information gathered during a trip [5, 6]. We study how incomplete link observations affect the feasibility of such constraints. Our model uses route disruption as the failure event, enabling exact risk evaluation in controlled experiments. Arrival-time reliability also depends on delay distributions and correlations, which lie beyond this model.

The estimator pools observations over several spatial scales. Larger neighborhoods reduce sampling uncertainty but increase the possible mismatch between pooled and target probabilities. We account for both effects before choosing a radius. A simultaneous concentration bound then supports subsequent route optimization without an additional data split. This matters because selecting a favorable route from many candidates can amplify errors in uncalibrated link estimates.

Our experiments test three distinct obstacles to transfer: missing individual links, a contiguous sensing gap, and rapid spatial variation near a hazard concentration. A piecewise-smooth field additionally tests transfer restricted by a known class boundary. We compare pooled confidence with a stronger single-source transfer baseline, rather than attributing every

improvement over direct observations to pooling itself. The evaluation reports route acceptance, true risk, conditional detour, and simultaneous model coverage separately.

Related Work

Shortest-path methods exploit additive costs, while stochastic routing introduces uncertain transitions, delays, or arrival events [7–9]. Transit formulations optimize on-time arrival using decisions at stations and uncertain service arrivals [10]. These models motivate distinguishing a route fixed before departure from a policy that responds to observations during travel. Transportation networks also impose computational requirements beyond a single query [11]. Our label search handles an additive distance objective and a transformed probability constraint exactly on the selected network. Its worst-case label count remains exponential.

Resource-constrained shortest paths provide the underlying optimization structure. Lagrangian relaxation treats an additional path constraint through penalized shortest-path subproblems [12]. Approximation schemes offer another approach when a controlled objective error is acceptable [13]. We retain nondominated distance–risk labels to obtain an exact solution for each estimated network.

Confidence-based model learning connects observations to planning statements [14, 15]. Robust dynamic programming similarly optimizes over plausible transition models [16, 17]. We use a scalar upper bound for each link, which makes the relationship between estimation error and route risk transparent. The theoretical contribution is a finite-data composition of bounded spatial transfer with post-selection routing.

Analogies between state–action pairs also support safe exploration with unknown stochastic dynamics [18]. That setting couples observation collection to the evolving policy. Our analysis conditions on a fixed observation design and then certifies every route selected from the resulting bounds. This restriction isolates the effect of spatial estimation from the difficulty of choosing safe exploratory actions.

Constrained planning requires tracking quantities that determine future feasibility. Consumption

models track a replenishable resource [19]; our search tracks accumulated route risk. The analogy concerns the planning state, while the constraint semantics and algorithms differ.

We restrict the cited literature to publications available by December 31, 2022. We conducted the computational study in September 2026 using synthetic disruptions on a street graph acquired that month.

Problem Formulation

Network and Disruption Model

Let $G = (V, E)$ be a finite directed graph with positive edge lengths ℓ_e . An origin–destination query is a pair (s, t) . A feasible path P connects s to t through directed edges. Its distance is $L(P) = \sum_{e \in P} \ell_e$. Positive lengths ensure that an optimal solution can be taken to have no cycles.

Each edge has an unknown disruption probability $p_e \in [0, 1]$. Let X_e denote the disruption indicator for a deployment attempt. We define route failure as $\bigcup_{e \in P} \{X_e = 1\}$. Under independence, its probability is

$$\rho(P; p) = 1 - \prod_{e \in P} (1 - p_e). \quad (1)$$

The true-model optimization problem is

$$\min_{P: s \rightsquigarrow t} L(P) \quad \text{subject to} \quad \rho(P; p) \leq \alpha. \quad (2)$$

The algorithm observes data about p and may reject a query when it cannot certify a route.

The disruption event can represent loss of route service or a severe delay that prevents completion. Mapping it to an arrival-time guarantee requires a separate condition on travel time when no disruption occurs. Our experiments evaluate (1) directly. They contain no timetable, waiting-time, or passenger-choice model.

Observation Model and Spatial Regularity

Let $z_e \in \mathbb{R}^2$ be the projected midpoint of edge e . For observed edge e_{jk} , the training data contain Bernoulli observations y_{jk} with mean $z_{e_{jk}}$. The counts may be zero. We condition on the counts and observed-edge set in the analysis.

Assumption

(Sampling and stationarity). *Conditional on the probability field and sampling design, all training*

observations are independent. The counts and the candidate pooling neighborhoods are chosen independently of their observed outcomes. Training and deployment use the same probability field.

Assumption

(Spatial probability bound). A known constant L_p satisfies $|p_e - p_j| \leq L_p \|z_e - z_j\|_2$ for all pairs used in pooling. Pools may be restricted to known classes, with no transfer across class boundaries. The disruption event has the same meaning on every pooled edge.

Assumption 2.2 formalizes transferable information rather than realized traffic dependence. It can hold for a synthetic smooth field even when all disruption events are independent. In a measured network, the bound would require domain evidence or independent calibration. Blindly replacing it with geometric proximity removes the basis of the guarantee. Spatial learning similarly separates bounds on probability differences from the samples used to estimate those probabilities [2].

Assumption

(Deployment independence). Conditional on the fixed probability field, disruption indicators along a selected route are independent.

This assumption is stronger than the requirements for a union-bound risk estimate. PROTRIP addresses correlated travel-time information directly [5]. We retain independence here to obtain the exact product form in (1). Section 7 describes the conservative alternative when dependence is unspecified.

Spatial Confidence Transfer

Pooled Probability Estimation

Fix a candidate neighborhood $N_r(e)$. Its pooled count, empirical mean, and sample-weighted distance are

$$N_{er} = \sum_{j \in N_r(e)} n_j, \quad (3)$$

$$\hat{p}_{er} = \frac{1}{N_{er}} \sum_{j \in N_r(e)} \sum_{k=1}^{n_j} Y_{jk}, \quad (4)$$

$$\bar{d}_{er} = \frac{1}{N_{er}} \sum_{j \in N_r(e)} n_j \|z_e - z_j\|_2. \quad (5)$$

These expressions are used only when $N_{er} > 0$. Let $\bar{p}_{er}$. Let p_j denote the correspondingly weighted average

of the true . Assumption 2.2 implies

$p_e \leq \bar{p}_{er} + L_p \bar{d}_{er}$. The pooled observations are independent but can have different means. Their average still admits a Bernoulli Chernoff bound. For any real λ , concavity of $\log(1 - p + pe^\lambda)$ in p gives

$$\mathbb{E} \exp(\lambda \sum Y_{jk}) \leq (1 - \bar{p}_{er} + \bar{p}_{er} e^\lambda)^{N_{er}}. \quad (7)$$

Optimizing the exponential bound yields the usual binary relative-entropy tail expression. This derivation is a scalar instance of concentration for bounded independent observations [20].

Define the binary relative entropy

$$\text{kl}(x \| q) = x \log \frac{x}{q} + (1 - x) \log \frac{1 - x}{1 - q}, \quad (8)$$

with the conventional endpoint limits. For a confidence allocation η , define

$$U(x, N, \eta) = \sup\{q \in [x, 1] : N \text{kl}(x \| q) \leq \log(1/\eta)\}. \quad (9)$$

We set when . For , a monotone bisection computes the endpoint. The same Bernoulli relative-entropy geometry underlies KL-UCB confidence indices for stochastic bandits [21]. Here, the error allocation covers a fixed collection of spatial pools; the coverage proof also accounts for unequal observation means.

Radius Selection and Simultaneous Coverage

Let K be the number of predetermined candidate neighborhoods and $M = |E|$. We allocate $\eta = \delta/(MK)$ to each edge–neighborhood pair. The transferred bound is

$$u_{er} = \min\{1, U(\hat{p}_{er}, N_{er}, \delta/(MK)) + L_p \bar{d}_{er}\}, \quad u_e = \min_r u_{er}. \quad (10)$$

A neighborhood with no observations receives bound one. The smallest neighborhood is the edge's own data. Larger neighborhoods contain eligible edge midpoints within a specified radius. The eligibility classes and radii are fixed before observing outcomes.

Proposition 3.1 (Simultaneous transfer coverage). Under Assumptions 2.1 and 2.2,

$$\mathbb{P}_D(p_e \leq u_e \text{ for every } e) \geq 1 - \delta. \quad (11)$$

The result allows selecting each radius using the observed values in (10).

Proof. For each nonempty pool, the Chernoff inversion gives

Equation (6) then bounds p_e by u_{er} on the same event. Empty pools are valid because probabilities

cannot exceed one. A union bound over all MK candidates makes every bound valid simultaneously. Their minimum is therefore also an upper bound. The pools can overlap because the union bound requires no independence between their coverage events. $\square$

Radius selection trades sampling uncertainty against transfer bias. A large pool can substantially increase N_{er} while also increasing $L_p \bar{u}_{er}$. The method retains the smallest certified upper bound among those alternatives. If the prescribed spatial bound is loose, more distant observations can remain unhelpful even with large sample counts. This mechanism explains persistent abstention in the experiments.

Risk-Constrained Routing

Additive Risk Formulation

Define the conservative edge weight

$$w_e = -\log(1 - u_e), \quad B_\alpha = -\log(1 - \alpha). \quad (11)$$

An edge with $u_e = 1$ has infinite risk weight. The certified routing problem becomes

$$\min_{P: \rho(P) \leq \alpha} \sum_{e \in P} w_e \quad \text{subject to} \quad \sum_{e \in P} w_e \leq B_\alpha. \quad (12)$$

The logarithmic transformation preserves the multiplicative reliability requirement.

Proposition

(Reliability of accepted routes). *Under Assumptions 2.1–2.3, with training-data probability at least $1 - \delta$, every accepted route satisfies $\rho(P; p) \leq \alpha$. The statement holds for any number of queries answered using the same simultaneous confidence bounds.*

Proof. On the event of Proposition 3.1, $p_e \leq u_e$ for all edges. The failure expression (1) is monotone in every edge probability. Any route satisfying (12) therefore has $\rho(P; p) \leq \rho(P; u) \leq \alpha$. The argument is simultaneous over all graph paths, including paths selected after inspecting the data. $\square$

The guarantee separates model confidence from deployment reliability. The value α bounds the chance that the training data produce an invalid simultaneous model. The value δ bounds route failure on the coverage event. For a rule that can abstain, an unconditional

joint bound is $|\mathbb{P}(\text{accepted} \cap \text{failure})| \leq \alpha \mathbb{P}(\text{accepted}) + \delta$. This does not provide the same conditional bound after conditioning on acceptance across datasets.

Nondominated Label Search

We solve (12) with a label search over accumulated distance and risk. A label at node v dominates when both coordinates are no larger and at least one is smaller. Dominated labels cannot produce a better feasible completion. A shortest-risk distance from each node to the goal provides an additional admissible pruning rule.

Algorithm

Spatial confidence routing

Positive edge lengths and nonnegative risk weights justify the queue ordering and dominance checks. The first feasible goal label has minimum length among certified paths. The number of nondominated labels can grow exponentially in the worst case. We therefore make no polynomial-time claim for the constrained search. The experiment uses the complete search without an arbitrary label cap.

Single-Source Transfer Baseline

Pooling should be compared with information transfer that does not aggregate samples. Let $d_j = U(\hat{p}_j, n_j, \delta/M)$ be a direct upper bound. The envelope estimator is

$$u_e^{\text{env}} = \min \left\{ 1, \min_{j \in \mathcal{C}(e)} [d_j + L_p \|z_e - z_j\|_2] \right\}, \quad (13)$$

where $\mathcal{C}(e)$ contains eligible links, including itself. On simultaneous direct coverage, every candidate is valid by the spatial bound. Selecting the smallest candidate therefore retains coverage with probability at least $1 - \delta$. The envelope can certify an unobserved edge from one well-observed neighbor. Pooling can improve it when several individually weak neighbors provide enough joint evidence to offset their bias allowances. Neither estimator uniformly dominates the other because their confidence allocations differ.

Experimental Design

Road Network and Routing Queries

We used OSMnx to acquire and project a directed street graph centered on Champaign, Illinois [22]. The frozen OpenStreetMap snapshot contains 345 nodes and 1031 directed links in its largest strongly




Figure 1: Synthetic hazards on the OpenStreetMap street graph. (a) Smooth background field. (b) Concentrated hazard with a 350 m spatial scale. (c) Piecewise-smooth field with a known class boundary. The circle and square mark one fixed origin-destination pair in every panel; the experiment uses 24 pairs. Coordinates are local projected distances. Road geometry: © OpenStreetMap contributors, ODbL 1.0.

connected component. Parallel arcs between the same ordered nodes were reduced to the shortest arc. The saved GraphML file and its hash fix the network used by every run. OpenStreetMap supplied geometry and connectivity; all disruption probabilities and observations were generated synthetically.

Trip selection used geometry before generating any disruption observations. We selected eight origin-destination pairs in each shortest-distance

stratum: 500–1200 m, 1200–2000 m, and 2000–3200 m. The resulting 24 pairs remained fixed across methods, fields, and datasets. This design includes short local trips and longer trips that cross several corridors. It avoids selecting queries after inspecting a method's acceptance decisions.

Disruption Probability Models

Let (x_e, y_e) be the projected edge midpoint after

Algorithm 1. Spatial confidence routing

Require: Graph G , observations D , radii R , L_p , δ , query (s, g) , threshold α

- 1: **for** each edge e and radius $r \in R$ **do**
- 2: Compute N_{er} , $\hat{p}_{er}$, and $\bar{d}_{er}$
- 3: Construct u_{er} by (10)
- 4: **end for**
- 5: Set $u_e \leftarrow \min_r u_{er}$ and $w_e \leftarrow -\log(1 - u_e)$
- 6: Compute minimum remaining risk $h(v)$ from each node to g
- 7: Initialize a distance-ordered queue with label $(s, 0, 0)$
- 8: **while** the queue is nonempty **do**
- 9: Pop the label (v, d, r) with smallest distance
- 10: **if** $v = g$ **then**
- 11: **return** the path represented by this label
- 12: **end if**
- 13: **for** each edge $e = (v, v')$ **do**
- 14: Form $(d', r') \leftarrow (d + \ell_e, r + w_e)$
- 15: **if** $r' + h(v') \leq B_\alpha$ and the new label is nondominated **then**
- 16: Insert the new label and remove labels it dominates
- 17: **end if**
- 18: **end for**
- 19: **end while**
- 20: **return** abstain




Figure 2: Route acceptance at a 10% risk limit. Columns show smooth, concentrated-hazard, and boundary fields. Panels (a)–(c) use complete sensing; panels (d)–(f) use scattered sensing with observation probability 0.65. Curves average 24 trip pairs within each dataset and then average across 24 datasets. Bands give 95% Student- intervals across datasets. Oracle acceptance marks true feasibility; higher acceptance by Kernel can include unsafe routes.

subtracting the coordinatewise minimum. The smooth field is

$$p_e^S = 0.0015 + 1.2 \times 10^{-6} x_e + 0.6 \times 10^{-6} y_e. \quad (14)$$

We supplied $L_p = 1.5 \times 10^{-6} \text{ m}^{-1}$, which bounds its gradient norm. The concentrated-hazard field adds a Gaussian hotspot:

$$p_e^H = p_e^S + 0.018 \exp\left(-\frac{\|z_e - z_c\|_2^2}{2(350 \text{ m})^2}\right), \quad (15)$$

where z_c is the coordinatewise median midpoint. Its supplied bound was $3.3 \times 10^{-5} \text{ m}^{-1}$. The Gaussian gradient has maximum magnitude $0.018e^{-1/2}/350$, so this bound also includes the background gradient.

The boundary field is

$$p_e^B = p_e^S + 0.0121 \{x_e > 0.52x_{\max}\} \quad (16)$$

A known binary class marks the two sides of that boundary. Pools remain inside a class and use the smooth-field Lipschitz bound. This construction separates a known discontinuity from an unknown smooth variation. Figure 1 displays all three fields on the same road geometry.

Observation Coverage and Baselines

Complete sensing observed every link. Scattered sensing observed each link independently with probability 0.65, with the observed set fixed across budgets within a dataset. Corridor sensing omitted all links whose midpoint satisfied $0.40x_{\max} < x_e < 0.58x_{\max}$. Every observed link received or 100, 400, 1600, or 6400 independent Bernoulli observations. We generated 24 datasets, sharing each dataset across comparison methods. Candidate radii were 0, 60, 150, 300, and 600 m.

Direct used simultaneous confidence bounds from each link's own observations. Envelope used (13). Spatial used the multiscale pooled bound in (10). Kernel used the empirical pooled fraction from the radius selected by Spatial, without a confidence or bias allowance. No bias retained the pooled concentration term while setting the spatial correction to zero. Oracle solved the routing problem with the generating probabilities. All methods used the same exact constrained search.

The main comparison used $\alpha = 0.10$ and $\delta = 0.05$




Figure 3: Calibration and missing-coverage effects. (a) Simultaneous link coverage in the concentrated-hazard field under scattered sensing. (b) Fraction of accepted routes violating the threshold, aggregated across fields at 1600 observations; error bars summarize 95% Student- intervals over dataset-level fractions. (c) Acceptance with a contiguous sensing gap at 6400 observations, where S, H, and B denote the three hazard fields. Panel (c) compares two confidence-based transfer estimators under the same missing region.

A threshold sensitivity study used $\alpha \in \{0.05, 0.10, 0.20\}$ at 1600 observations under scattered sensing. Across fields, patterns, thresholds, and methods, the saved experiment contains 145,152 route decisions. This count includes repeated evaluation of the deterministic oracle for aligned comparisons. The

Table 1: Comparison at 1600 observations per sensed link and $\alpha=0.10$. Acceptance uses all queries. Detour averages accepted, truly feasible routes under scattered sensing. Violations count accepted routes whose exact risk exceeds the limit. Coverage is the fraction of training datasets whose entire upper-bound field covers truth; point estimates and the oracle have no estimated confidence statement.

Field	Method	Acceptance (%)		Scattered sensing		
		Complete	Scattered	Detour (%)	Violations	Coverage (%)
Smooth	Direct	11.5	1.4	0.06	0	100.0
	Envelope	58.7	58.2	0.01	0	100.0
	Spatial	91.7	91.7	0.24	0	100.0
	Kernel	100.0	100.0	0.00	5	—
	No bias	100.0	100.0	0.27	0	100.0
	Oracle	100.0	100.0	0.00	0	—
Hotspot	Direct	7.6	1.2	8.42	0	100.0
	Envelope	9.0	8.3	0.00	0	100.0
	Spatial	12.7	12.5	0.00	0	100.0
	Kernel	62.2	62.3	0.00	59	—
	No bias	42.5	41.8	0.02	0	0.0
	Oracle	58.3	58.3	0.00	0	—
Boundary	Direct	8.5	1.2	8.42	0	100.0
	Envelope	24.8	24.8	0.02	0	100.0
	Spatial	33.3	33.3	0.00	0	100.0
	Kernel	41.7	41.7	0.00	16	—
	No bias	33.7	33.3	0.00	0	95.8
	Oracle	41.7	41.7	0.00	0	—

independent units for uncertainty estimates are the 24 training datasets.

Evaluation Metrics and Validation

Acceptance is the fraction of all 24 queries for which a method returns a route. Oracle acceptance indicates the fraction that is feasible under the generating field. A point-estimate method can exceed oracle acceptance by returning an infeasible route. We therefore report the true violation count beside acceptance. Each accepted route's actual risk was computed exactly from (1).

Detour is route length divided by the corresponding oracle's constrained-optimal length, minus one. The table averages detour only over accepted, truly feasible routes. This conditioning prevents an infeasible shortcut from appearing as a routing improvement. It also means that detour summaries compare different accepted subsets when methods abstain differently.

Simultaneous coverage records whether every estimated upper bound covers its generating link probability within one dataset. This diagnostic is stricter than checking only the links selected by the planner. Uncertainty bands first average queries within a dataset, then use Student- intervals across datasets. Treating all query rows as independent would underestimate uncertainty because they share a learned network model.

The implementation checked constrained search against exhaustive simple-path enumeration on small directed graphs. A separate finite-outcome calculation checked the pooled confidence construction with unequal Bernoulli means. Additional audits recomputed risks and route lengths from saved path snapshots. These checks target the two operations that connect the theorem to the numerical results.

Results

Route Acceptance

Under scattered sensing with 1600 observations, Spatial accepted 91.7% of queries in the smooth field. Envelope accepted 58.2%. Both methods transferred information into unobserved links, so the difference measures the benefit of aggregating several sources under a bias allowance. Figure 2 shows that direct

observations catch up under complete sensing as data increase. Missing links continue to limit Direct under scattered sensing.

The concentrated hazard substantially reduced the benefit of transfer. Spatial accepted 12.5% of queries at the same budget and observation pattern. Its larger valid Lipschitz constant increased the bias allowance even for short transfers. The confidence estimator responded by using more local evidence and abstaining more often. This behavior follows from the stated model, rather than a numerical failure of the routing search.

Confidence Coverage and Risk Violations

Spatial produced 0 route-risk violations over the full experiment. Direct and Envelope also produced no violations. These observations are consistent with their conservative confidence constructions; they do not establish a zero failure probability outside the specified model. The theoretical statement remains a $1-\delta$ training-data guarantee.

Kernel produced 1569 violating route decisions, and No bias produced 264. A point estimate can rank routes effectively while still underestimating a risk near the acceptance threshold. The bias ablation additionally pools across a changing probability field without accounting for that change. As sampling uncertainty shrinks, the omitted spatial discrepancy can dominate its remaining error. Figure 3 compares simultaneous coverage with the observed routing consequences.

Sensing Gaps and Route Detours

Scattered gaps can leave observed links close enough for useful transfer. A missing corridor increases the distance to available evidence, leaving a positive bias allowance even when observation counts grow. Known class boundaries impose the same limitation by excluding otherwise nearby links from the pool. These cases require observation coverage to be assessed alongside sample count.

Many accepted queries had a shortest-distance route comfortably below the risk threshold, so certification added little conditional detour. Near the threshold, conservative estimates could select a longer route or reject the query. Acceptance and detour must be interpreted together because detour

is averaged only over accepted, feasible routes.

Extensions and Limitations

Correlated Disruptions and Arrival-Time Reliability

The product risk in (1) requires independent deployment disruptions. Spatially similar probabilities do not imply this independence. A shared incident can correlate several links even when their marginal probabilities satisfy a smoothness bound. When only marginals are known, the conservative condition

$$\sum_{e \in P} u_e < \alpha \quad (17)$$

controls route failure by the union bound without a dependence model. It changes the risk weights and generally increases abstention. We did not evaluate correlated incidents in the reported benchmark.

Arrival-time reliability needs an additional model of travel and waiting times. A binary disruption event can represent failure of a service guarantee, but its complement does not specify a complete arrival-time distribution. Stochastic transit policies and reliability-guaranteed routing address richer travel uncertainty and policy decisions [5, 6, 10]. Applying spatial transfer there would require bounds on transition or travel-time distributions and a dependence-aware planning recursion. The present result establishes the simpler marginal-estimation step.

Spatial Model Calibration and Adaptive Sampling

The experiments supply correct class labels and an analytic Lipschitz bound. Estimating either quantity from the same observations without correction would alter the coverage argument. A practical deployment could calibrate similarity on independent data, use a conservative physical bound, or introduce confidence sets for the spatial relation itself. The corresponding uncertainty must enter the route certificate.

The fixed-design concentration argument also excludes outcome-dependent stopping of training collection. Online route data are naturally adaptive because earlier estimates influence which links are visited next. Confidence sequences provide bounds that remain valid under continuous monitoring and data-dependent stopping [23]. Using them here would require a valid sequential observation model and confidence allocation across the eligible pools. The distinction matters when translating the finite-data

experiment into an operational learning system.

Benchmark Scope and Computational Complexity

The benchmark uses one urban street graph with several hazard and sensing configurations. Its geometry is real, but the disruption model is synthetic and does not fit traffic records. The experiment supports conclusions about estimator behavior under controlled assumptions. A transport study would additionally need multiple networks, observed incident data, and a suitable service or arrival event.

The confidence calculation can use sparse neighborhood matrices on larger graphs. The constrained path search can still accumulate many nondominated labels, especially when distance and risk strongly conflict. Restricted-shortest-path approximation schemes suggest alternatives to complete label enumeration [13]. Any approximate implementation would need to retain certified feasibility and quantify its objective error. The current implementation retains the exact search and saves the graph snapshot, observations, and route-level outcomes.

Conclusion

Bounded spatial variation can make incomplete observations useful for reliable routing. Pooling several sources improved route certification beyond single-source transfer when the spatial bias allowance was small. Concentrated hazards and contiguous sensing gaps reduced that advantage. The experiments also showed why route acceptance must be evaluated alongside true violations and simultaneous model coverage. Extending the result to transit arrival guarantees will require richer travel-time and dependence models.

References

1. Melkior Ornik, Jie Fu, Niklas T. Lauffer, W. K. Perera, Mohammed Alshiekh, Masahiro Ono, and Ufuk Topcu. Expedited Learning in MDPs with Side Information. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 1941–1948, 2018. doi: 10.1109/cdc.2018.8619134.
2. Pranay Thangeda and Melkior Ornik. Safety-Guaranteed, Accelerated Learning in MDPs with Local Side Information. In *2020 American Control Conference (ACC)*, pages 1099–1104, 2020. doi: 10.23919/acc45564.2020.9147372.
3. Pranay Thangeda. Efficient learning and planning using spatial side information. Master's thesis, University of

- Illinois at Urbana-Champaign, 2020. M.S. thesis.
4. Pranay Thangeda, Melkior Ornik, and Ufuk Topcu. Expedited Online Learning With Spatial Side Information. *IEEE Transactions on Automatic Control*, 2022. doi: 10.1109/tac.2022.3153278. Early access.
5. Pranay Thangeda and Melkior Ornik. PROTRIP: Probabilistic Risk-Aware Optimal Transit Planner. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6, 2020. doi: 10.1109/itsc45102.2020.9294500.
6. Wanzheng Zheng, Pranay Thangeda, Yagiz Savas, and Melkior Ornik. Optimal Routing in Stochastic Networks with Reliability Guarantees. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 3521–3526, 2021. doi: 10.1109/itsc48978.2021.9564444.
7. E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959. doi: 10.1007/bf01386390.
8. Dimitri P. Bertsekas and John N. Tsitsiklis. An Analysis of Stochastic Shortest Path Problems. *Mathematics of Operations Research*, 16(3):580–595, 1991. doi: 10.1287/moor.16.3.580.
9. Yu (Marco) Nie and Xing Wu. Shortest path problem considering on-time arrival probability. *Transportation Research Part B: Methodological*, 43(6):597–613, 2009. doi: 10.1016/j.trb.2009.01.008.
10. Yang Liu, Sebastien Blandin, and Samitha Samaranyake. Stochastic on-time arrival problem in transit networks. *Transportation Research Part B: Methodological*, 119:122–138, 2019. doi: 10.1016/j.trb.2018.11.013.
11. Hannah Bast, Daniel Delling, Andrew Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F. Werneck. Route Planning in Transportation Networks. In *Algorithm Engineering*, pages 19–80, 2016. doi: 10.1007/978-3-319-49487-6_2.
12. Gabriel Y. Handler and Israel Zang. A dual algorithm for the constrained shortest path problem. *Networks*, 10(4):293–309, 1980. doi: 10.1002/net.3230100403.
13. Refael Hassin. Approximation Schemes for the Restricted Shortest Path Problem. *Mathematics of Operations Research*, 17(1):36–42, 1992. doi: 10.1287/moor.17.1.36.
14. Michael Kearns and Satinder Singh. Near-Optimal Reinforcement Learning in Polynomial Time. *Machine Learning*, 49(2-3):209–232, 2002. doi: 10.1023/a:1017984413808.
15. Alexander L. Strehl and Michael L. Littman. An analysis of model-based Interval Estimation for Markov Decision Processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008. doi: 10.1016/j.jcss.2007.08.009.
16. Arnab Nilim and Laurent El Ghaoui. Robust Control of Markov Decision Processes with Uncertain Transition Matrices. *Operations Research*, 53(5):780–798, 2005. doi: 10.1287/opre.1050.0216.
17. Garud N. Iyengar. Robust Dynamic Programming. *Mathematics of Operations Research*, 30(2):257–280, 2005. doi: 10.1287/moor.1040.0129.
18. Melrose Roderick, Vaishnavh Nagarajan, and J. Zico Kolter. Provably Safe PAC-MDP Exploration Using Analogies. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130, pages 1216–1224. PMLR, 2021. URL <https://proceedings.mlr.press/v130/roderick21a.html>.
19. František Blahoudek, Tomáš Brázdil, Petr Novotný, Melkior Ornik, Pranay Thangeda, and Ufuk Topcu. Qualitative Controller Synthesis for Consumption Markov Decision Processes. In *Computer Aided Verification*, pages 421–447, 2020. doi: 10.1007/978-3-030-53291-8_22.
20. Wassily Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830.
21. Aurélien Garivier and Olivier Cappé. The KL-UCB Algorithm for Bounded Stochastic Bandits and Beyond. In *Proceedings of the 24th Annual Conference on Learning Theory*, volume 19, pages 359–376. PMLR, 2011. URL <https://proceedings.mlr.press/v19/garivier11a.html>.
22. Geoff Boeing. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017. doi: 10.1016/j.compenvurbsys.2017.05.004.
23. Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. doi: 10.1214/20-aos1991.