Bayesian Experimental Design for Autonomous Vehicle Testing

Lawrence R. De Geest
Zoox
Foster City, CA, USA
ldegeest@zoox.com

ABSTRACT

Autonomous vehicle development requires collecting large amounts of on-road mileage to train and evaluate driving systems across a heterogeneous road network. However, allocating mileage efficiently is challenging because driving environments vary substantially and each additional traversal is costly. We frame mileage collection as a sequential experimental design problem in which a fixed budget of traversals must be allocated across roads to efficiently learn heterogeneous road-level event rates. We present an empirical Bayes approach that pools information across roads while adaptively directing additional mileage toward locations where further observations are most informative. We explore the behavior of the proposed allocation strategy using simulated road networks with heterogeneous latent event rates.

1 INTRODUCTION

Autonomous driving systems require extensive on-road testing to characterize performance across a large and heterogeneous operating domain. Yet collecting this experience is costly, and many outcomes of interest occur too infrequently to measure precisely everywhere. Collisions and other safety-critical events provide an extreme example, but the same problem arises whenever important software behaviors or failure modes are sparse relative to the amount of driving required to observe them. Consequently, data collection is not simply a question of accumulating mileage, but of deciding *where that mileage should be collected.*

This distinction matters because roads are not equally informative. Road geometry, intersections, traffic composition, and other characteristics generate different driving scenarios and underlying levels of risk [1, 4, 5]. The distribution of driving exposure itself can also vary substantially across space and time, making where mileage is accumulated consequential for performance evaluation [2]. Figure 1 illustrates this heterogeneity using publicly available human collision data from San Francisco, showing that both raw collisions and estimated baseline collision risk vary considerably across the city.[1] An additional traversal of a frequently driven, well-understood road therefore need not provide as much information as a traversal of a poorly understood road. A mileage collection policy that ignores this heterogeneity can accumulate substantial exposure while learning relatively little about important portions of the road network.

[1]Data from https://data.sfgov.org/Public-Safety/Traffic-Crashes-Resulting-in-Injury/ubvf-ztfx/about_data

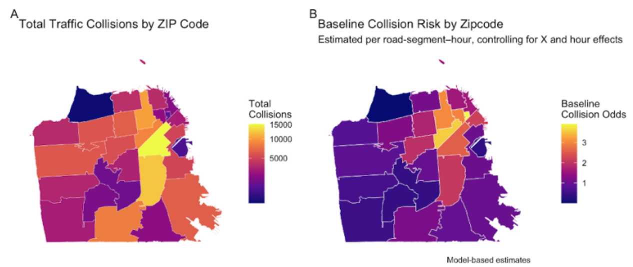

Figure 1: Spatial heterogeneity in human collision risk in San Francisco. Panel A shows observed collisions by ZIP code; Panel B shows estimated baseline collision risk after controlling for observable road-segment and temporal characteristics.

We frame mileage collection as a sequential experimental design problem. Consider a road network containing many road segments, each associated with an unknown latent performance or risk. Traversing a segment provides another exposure and a set of observable outcomes. Given a finite collection budget, the experimenter's problem is to allocate traversals across the network to learn this latent surface as efficiently as possible. Unlike a conventional randomized experiment, each observation changes what we know about the experimental units and therefore the value of collecting the next observation.

This Bayesian formulation naturally turns estimation into experimental design. Given our current beliefs about each road, we can value another traversal by its expected information gain and allocate exposure toward roads where we have the most to learn. Because information exhibits diminishing returns, the value of additional traversals falls as a road becomes better understood. These values then serve as edge weights on a road-network graph. Experimental design determines where additional mileage is most informative, while a downstream routing problem determines how that mileage can feasibly be collected.

In this paper, we develop an empirical Bayes approach to this problem and study its behavior in simulated road networks. We use sequential posterior updating to value additional traversals and compare adaptive mileage allocation with alternative collection strategies. Our broader objective is to show how the practical question of *where should an autonomous vehicle drive next?* can be viewed as a Bayesian optimal experimental design problem.

2 LEARNING ROAD-LEVEL EVENT RATES

Let Y_i denote the number of observed events on road segment i and N_i the number of traversals of that segment. We assume that events arrive according to a road-specific latent event rate λ_i,

$$Y_i \mid \lambda_i, N_i \sim \text{Poisson}(N_i \lambda_i). \tag{1}$$

Road-level event rates may vary systematically with observable characteristics such as road geometry, intersection type, speed limit, traffic volume, and the surrounding environment. A substantial transportation literature documents both systematic relationships between these characteristics and crash frequency and residual spatial and unobserved heterogeneity across road segments [1, 4, 5]. For example, adjacent road segments exhibit correlated crash rates [5], while the effects of characteristics such as speed limits, number of lanes, intersections, and traffic exposure have been shown to vary spatially [1]. We therefore model the expected event rate for road i as

$$\log \bar{\lambda}_i = X_i' \beta + u_{g(i)}, \tag{2}$$

where X_i contains observable road characteristics and $u_{g(i)}$ is a group-level geographic random effect. The geographic grouping could represent a ZIP code, neighborhood, or other collection of nearby road segments intended to capture shared latent environmental factors.

To allow individual roads to deviate from this prediction, we place a Gamma distribution around the expected rate,

$$\lambda_i \sim \text{Gamma}\left(\alpha, \frac{\alpha}{\bar{\lambda}_i}\right), \tag{3}$$

where we use the shape–rate parameterization. Integrating over λ_i produces a negative binomial distribution for Y_i, allowing event counts to be overdispersed relative to a Poisson model. Negative binomial models and extensions that explicitly accommodate unobserved heterogeneity are commonly used for road-segment crash counts [1]. We estimate the regression, geographic effects, and dispersion parameter using historical road-level data.

The Gamma–Poisson representation also provides a convenient empirical Bayes update for each road. Conditional on the estimated model, observing Y_i events over N_i traversals gives

$$\lambda_i \mid Y_i, N_i \sim \text{Gamma}\left(\alpha + Y_i, \frac{\alpha}{\bar{\lambda}_i} + N_i\right). \tag{4}$$

The corresponding posterior mean is

$$\mathbb{E}[\lambda_i \mid Y_i, N_i] = \frac{\alpha + Y_i}{\alpha / \bar{\lambda}_i + N_i}. \tag{5}$$

This estimator partially pools information across the road network. Roads with little exposure are shrunk toward the event rate predicted by observable characteristics and their surrounding geography, while roads with substantial exposure are increasingly determined by their own observed outcomes.

3 ALLOCATING EXPOSURES: BAYESIAN EXPERIMENTAL DESIGN

Crucially for experimental design, the posterior captures not only our estimate of each road's event rate but also our uncertainty about that estimate. We use this posterior distribution to determine where additional exposure is most informative.

Suppose mileage collection occurs in discrete planning periods, indexed by t. At the beginning of each period, the planner receives a budget of B_t traversals that must be allocated across I road segments. Let $n_{it} \in \mathbb{Z}_+$ denote the number of additional traversals assigned to road i during period t. The allocation must satisfy

$$\sum_{i=1}^{I} n_{it} \leq B_t. \tag{6}$$

At the beginning of period t, the current posterior distribution for the event rate on road i is

$$\lambda_i \mid \mathcal{D}_t \sim \text{Gamma}(a_{it}, b_{it}), \tag{7}$$

where $\mathcal{D}_t$ denotes all observations collected prior to period t. If road i is traversed once more and $Y_{i,t+1}$ events are observed, conjugacy gives

$$\lambda_i \mid \mathcal{D}_t, Y_{i,t+1} \sim \text{Gamma}\left(a_{it} + Y_{i,t+1}, b_{it} + 1\right). \tag{8}$$

Because the outcome of the next traversal is not known when the allocation decision is made, we evaluate a traversal by its expected information gain. For a single additional traversal, define

$$V_{it} = \mathbb{E}_{Y_{i,t+1} \mid \mathcal{D}_t}\left[D_{\text{KL}}\left(p(\lambda_i \mid \mathcal{D}_t, Y_{i,t+1}) \| p(\lambda_i \mid \mathcal{D}_t)\right)\right]. \tag{9}$$

Under the Gamma–Poisson model, the posterior predictive distribution of the outcome from one additional traversal is Negative Binomial,

$$Y_{i,t+1} \mid \mathcal{D}_t \sim \text{NegBin}\left(a_{it}, \frac{b_{it}}{b_{it}+1}\right). \tag{10}$$

Thus the expected information gain can be written explicitly as

$$V_{it} = \sum_{y=0}^{\infty} p(y \mid a_{it}, b_{it}) D_{\text{KL}}[\text{Gamma}(a_{it} + y, b_{it} + 1) \| \text{Gamma}(a_{it}, b_{it})]. \tag{11}$$

More generally, let $V_{it}(n)$ denote the expected information gained from allocating n additional traversals to road i during period t. The planner's weekly experimental design problem is

$$\max_{\{n_{it}\}} \sum_{i=1}^{I} V_{it}(n_{it}) \tag{12}$$

subject to

$$\sum_{i=1}^{I} n_{it} \leq B_t, \qquad n_{it} \in \mathbb{Z}_+. \tag{13}$$

The marginal information value of the kth traversal on road i is

$$\Delta V_{it}(k) = V_{it}(k) - V_{it}(k-1). \tag{14}$$

As additional observations accumulate, posterior uncertainty generally decreases and the marginal information value of further exposure declines. This produces a natural exploration mechanism: roads that have already been extensively sampled become relatively less valuable, while poorly understood roads receive greater weight.

A simple implementation therefore allocates traversals sequentially within each planning period. Starting from the current posterior distributions, the next traversal is assigned to the road with the greatest marginal expected information gain,

$$i^* = \arg\max_i \Delta V_{it}(n_{it} + 1), \tag{15}$$

and the process continues until the weekly traversal budget is exhausted. At the end of the period, observed event counts are incorporated into the posterior distributions and the allocation problem is solved again for the next period. While more frequent online updating is possible, in practice mileage collection is often organized over fixed planning horizons because changing routes and collection priorities entails nontrivial operational coordination and logistical constraints.

4 SIMULATION

We evaluate the proposed allocation rule using simulated event outcomes on a road network based on San Francisco road geometry. Road-level event rates are generated from the Poisson–Gamma model in Section 2, with observable road characteristics and geographic effects determining the prior mean and additional Gamma heterogeneity across segments. We assume historical exposure is uneven across roads, as a fixed fleet serving competing operational and data-collection priorities naturally accumulates more mileage on some parts of the network than others. We estimate the empirical Bayes model from these historical observations and then simulate 20 weekly planning periods with a fixed budget of 3,000 traversals per week.

We compare three allocation policies: uniform allocation, a business-as-usual policy that allocates mileage in proportion to historical exposure, and the proposed greedy Bayesian policy that assigns exposure according to marginal expected information gain. At the beginning of each week, the Bayesian policy constructs the week's allocation from the current posterior distributions. Realized event counts are then observed, posteriors are updated, and the allocation problem is re-solved the following week. For computational efficiency, the greedy rule is implemented in blocks of 10 traversals. All policies face the same simulated road-level environment and common random shocks within each Monte Carlo replication.

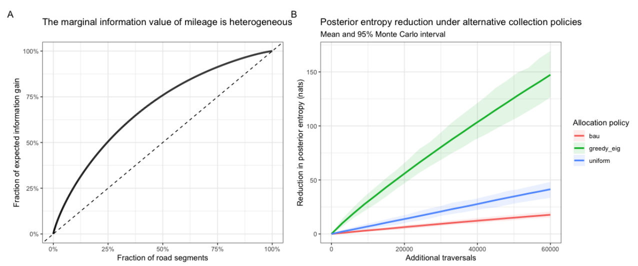

Figure 2: Simulation results. Panel A shows the concentration of marginal expected information across road segments; Panel B shows posterior entropy reduction under alternative mileage-allocation policies.

Figure 2 illustrates the source and magnitude of the gains from adaptive allocation. Marginal information value is concentrated across the road network, with a relatively small share of segments accounting for a disproportionate share of potential information gain. By directing traversals toward these segments and adapting as posterior beliefs evolve, the Bayesian policy reduces posterior entropy substantially faster than either uniform or business-as-usual allocation under the same mileage budget.

5 FROM EXPOSURE ALLOCATION TO ROUTING

The allocation problem above determines *where* additional mileage is most informative but does not determine how a vehicle should physically collect that mileage. To connect the experimental design to routing, represent the road network as a graph $G=(V,E)$, where each road segment $i \in E$ is associated with an information value function $V_{it}(\cdot)$. A downstream routing problem then selects a feasible set of road traversals that maximizes the information collected subject to operational constraints,

$$\max_{\mathbf{x}_t \in \mathcal{R}(B_t)} \sum_{i \in E} V_{it}(x_{it}), \tag{16}$$

where x_{it} denotes the number of times road segment i is traversed, $V_{it}(x_{it})$ is the total expected information gained from those traversals, and $\mathcal{R}(B_t)$ denotes the set of routes satisfying network connectivity and the available mileage or time budget. Equivalently, the kth traversal of an edge carries marginal value $\Delta V_{it}(k)$, allowing the routing objective to reflect the diminishing information value of repeated exposure. Thus, Bayesian experimental design supplies the value of traversing each edge, while the routing layer determines how to collect those exposures feasibly. As planned exposure accumulates, the information weights are updated to reflect diminishing marginal value.

6 DISCUSSION

The proposed framework provides a practical way to convert heterogeneous uncertainty across a road network into explicit values for additional data collection. By combining partial pooling with sequential experimental design, it directs scarce mileage toward locations where additional observations are expected to be most informative rather than simply accumulating exposure uniformly. The resulting information values can also be passed to a downstream routing problem, separating the statistical question of *where to learn* from the operational question of *how to collect the mileage*.

Several limitations remain. First, the quality of the allocation depends on the event-rate model and, where surrogate outcomes are used, on the stability of their relationship with the underlying target of interest [3]. The empirical Bayes model should therefore be periodically re-estimated as additional mileage accumulates and as changes in the driving system, operating domain, or road environment alter the relationship between observable road characteristics and event rates. Second, our weekly adaptive policy optimizes each planning period independently and does not account explicitly for the continuation value of current observations in future periods. Finally, the present analysis treats exposure allocation separately from routing. Incorporating network feasibility, traversal costs, and other operational constraints directly into the experimental design problem is a natural direction for future work.

More broadly, this framework illustrates how tools from sequential experimentation can be used to allocate scarce physical assets operating in uncertain environments. The same basic problem arises whenever deploying an asset both accomplishes an operational task and generates information that improves future deployment decisions. Dynamic programming can account for the continuation value of information, while bandit methods provide alternative approaches to balancing exploration of poorly understood environments against exploitation of known high-value locations. Future work could also extend the framework using tools from optimal stopping to jointly determine where to collect additional mileage and when continued sampling of a road or region no longer justifies its opportunity cost. Together, these approaches provide a broader toolkit for treating physical asset deployment as an adaptive experimental-design problem rather than a fixed data-collection plan.

REFERENCES

[1] Eskindir Ayele Atumo, Haibo Li, and Xinguo Jiang. 2023. Segment-Level Spatial Heterogeneity of Arterial Crash Frequency Using Locally Weighted Generalized Linear Models. *Transportation Research Record* 2677, 3 (2023), 1637–1653. https://doi.org/10.1177/03611981221126510

[2] Yi Chen, John M. Scanlon, Kristofer D. Kusano, Timothy McMurry, and Trent Victor. 2025. Dynamic Benchmarks: Spatial and Temporal Alignment for ADS Performance Evaluation. *Transportation Research Record: Journal of the Transportation Research Board* (2025). https://doi.org/10.1177/03611981251398744

[3] Yechen Li, Shantanu Shahane, Shoshana Vasserman, Carolina Osorio, Yi-fan Chen, Ivan Kuznetsov, Kristin White, Justyna Swiatkowska, Neha Arora, and Feng Guo. 2026. From Lagging to Leading: Validating Hard Braking Events as High-Density Indicators of Segment Crash Risk. *arXiv preprint arXiv:2601.06327* (2026). https://doi.org/10.48550/arXiv.2601.06327

[4] Shi Qiu, Hanzhang Ge, Zheng Li, Zhixiang Gao, and Chengbo Ai. 2024. Network-level crash risk analysis using large-scale geometry features. *Accident Analysis & Prevention* 207 (2024), 107746. https://doi.org/10.1016/j.aap.2024.107746

[5] Qiang Zeng, Huiying Wen, Helai Huang, and Mohamed Abdel-Aty. 2017. A Bayesian spatial random parameters Tobit model for analyzing crash rates on roadway segments. *Accident Analysis & Prevention* 100 (2017), 37–43. https://doi.org/10.1016/j.aap.2016.12.023