Journal of Urban Development and Smart Cities
2025, 2(1), 183-195

ADAPTIVE SURVEY DESIGN FOR SMART-CITY GEOTECHNICAL MAPPING IN SETAGAYA, TOKYO

A. M. Affif
Alan March
Yulesta Putra

Accurate prediction of bearing-layer depth is fundamental to resilient urban development, yet static subsurface maps offer limited guidance once additional site investigations must be prioritized under budget constraints. This study presents an adaptive geotechnical decision framework for Setagaya, Tokyo, using the published 433-site geotechnical data regime reported for the ward. The empirical basis comprises a 58.1 km^2 study area, a data density of 7.46 observations per km^2, and the core predictors used in the published benchmark—latitude, longitude, elevation, and bearing-layer depth.[1] Because the point-level coordinates are not publicly reproduced, the analysis is implemented as a constrained simulation that preserves the reported case definitions, model settings, and benchmark error scale. A hybrid residual bagging-kriging architecture is developed in which bagged trees estimate the nonlinear mean response and ordinary kriging interpolates the structured residual field. Predictive uncertainty is quantified by combining bootstrap dispersion and kriging variance, and a sequential acquisition rule allocates new boreholes by jointly maximizing uncertainty, predicted depth gradient, and local data sparsity. Relative to the published Setagaya bagging benchmark (RMSE 1.34 m; MAE 0.86 m), the proposed framework reduces RMSE to 1.08 m and MAE to 0.69 m while achieving 90% interval coverage of 0.89. Under a fixed drilling budget, uncertainty-aware allocation decreases holdout error 20–23% faster than random densification after twenty additional borings. The resulting workflow aligns geotechnical analytics with the needs of smart-city planning by linking prediction, uncertainty assessment, and field deployment inside a single GeoICT-oriented decision loop.

Index Terms — smart cities; geotechnical investigation; bagging; kriging; active learning; adaptive sampling; GeoICT; urban resilience

https://doi.org/10.66033/judsc2025-219

INTRODUCTION

Smart-city infrastructure depends not only on the predictive accuracy of analytical models, but also on the quality, timing, and spatial relevance of the decisions those models enable. In contemporary urban governance, digital systems are increasingly expected to do more than describe existing conditions; they must support resource allocation, reduce uncertainty in planning, and guide interventions under budgetary and operational constraints [1, 2, 3, 4, 5]. This expectation is especially important in domains where physical infrastructure decisions are costly, spatially uneven, and difficult to reverse once implemented. Urban geotechnics is one such domain. The depth of the bearing layer strongly influences foundation design, excavation strategy, structural feasibility, construction cost, and exposure to settlement or subsidence-related risk. As a result, subsurface uncertainty is not merely a technical inconvenience: it is a planning problem with direct implications for safety, affordability, project scheduling, and long-term urban resilience [1, 6, 7, 8].

In many built-up cities, however, geotechnical information is acquired incrementally rather than comprehensively. Site investigations are often commissioned on a parcel-by-parcel basis, driven by immediate development demand rather than by a coordinated metropolitan survey strategy. This creates an uneven data landscape in which some neighbourhoods become relatively well characterised while others remain sparsely observed. Under such circumstances, municipalities and private developers require more than a static predictive map of expected subsurface conditions. They require a rational framework for determining where the next investigation should be conducted, how uncertainty should be prioritised, and whether additional drilling will materially improve the quality of future decisions [2, 3, 5, 9]. In other words, the problem is not only one of interpolation, but also one of sequential urban information design.

Setagaya, Tokyo provides an appropriate empirical setting for examining this problem. The published geotechnical benchmark for the ward used 433 investigations to predict bearing-layer depth and demonstrated a pronounced accuracy advantage for bagging over ordinary kriging, with reported RMSE values of 1.34 m and 3.56 m, respectively [1]. The same benchmark further showed that kriging error was lower in locations surrounded by denser nearby observations, reporting a correlation coefficient of $r = -0.62$ between local data abundance and prediction error [1]. These results establish two important points. First, ensemble learning can substantially improve predictive performance in heterogeneous urban subsurface settings. Second, spatial information density still matters, even when strong predictive models are available. Taken together, these findings point to an unresolved smart-city and urban-analytics question: once a city already possesses an incomplete but informative subsurface model, how should geotechnical surveys be densified in a way that is both statistically efficient and operationally useful [1, 4, 5, 10]?

That question is directly relevant to the scope of research on urban development and smart cities because it connects spatial analytics, infrastructure resilience, data-informed governance, and implementation-oriented planning [2, 3, 4, 6, 7]. A geotechnical model that predicts well but cannot guide subsequent observation remains descriptively valuable, yet strategically incomplete. It can inform isolated site decisions, but it does not fully support coordinated urban management. By contrast, a system that estimates the subsurface, identifies where uncertainty remains high, and recommends where new borings are likely to generate the greatest informational gain functions as actionable smart-city infrastructure [3, 5, 8, 9, 11]. Such a system can support permitting decisions, reduce avoidable overdesign or underdesign, improve land-development sequencing, and strengthen resilience-oriented planning in areas vulnerable to geotechnical risk.

From a methodological standpoint, this problem sits at the intersection of three research traditions. The first is predictive urban analytics, which emphasizes the use of machine learning and spatial data science to model complex city systems [2, 4, 5]. The second is geostatistical reasoning, which remains essential for capturing local spatial dependence, residual structure, and uncertainty in geographically distributed phenomena [1, 6, 10]. The third is sequential or adaptive design, in which data collection is treated as an optimisation problem rather than a one-time exercise [3, 7, 8, 9]. Each of these traditions offers part of the solution, but none is sufficient on its own. Purely predictive approaches may deliver strong point estimates without an adequate uncertainty representation. Purely geostatistical approaches may preserve spatial structure but underperform in complex nonlinear settings. Non-adaptive survey strategies, meanwhile, may consume scarce investigation budgets without maximising information gain. A stronger framework therefore requires integration rather than substitution [1, 3, 5, 7, 10, 11].

This study addresses that need through a unified analytical framework with three linked objectives. The first objective is to retain the predictive strength of ensemble learning while recovering local spatial structure through residual kriging. This responds to the practical observation that urban subsurface variation often contains both nonlinear relationships and geographically structured residual behaviour. The second objective is to calibrate an uncertainty surface that is sufficiently informative to support closed-loop survey allocation. In a smart-city context, uncertainty is not a secondary output; it is a decision variable that should shape how new information is acquired and where public or private investigation resources are deployed [3, 4, 8, 9]. The third objective is to evaluate whether uncertainty-aware borehole placement reduces prediction error more efficiently than non-adaptive alternatives under a fixed investigation budget. This objective is especially important because urban planning decisions are almost always constrained by cost, time, and implementation capacity. Demonstrating better performance is useful, but demonstrating better performance per additional investigation is far more relevant to operational planning [2, 5, 7, 11].

The contribution of the study is therefore both methodological and operational. Methodologically, the paper integrates ensemble learning, geostatistics, and sequential design into a single workflow for urban subsurface prediction and adaptive survey planning [1, 3, 6, 10]. Operationally, it translates subsurface prediction into an implementable drilling policy that can support urban planning, permitting, redevelopment control, and resilience-oriented infrastructure management [2, 4, 5, 8, 9, 11]. By moving from static prediction toward uncertainty-aware action, the study reframes geotechnical mapping as part of the broader smart-city agenda: not simply knowing the city better, but deciding more intelligently about where to learn next.

PUBLISHED EMPIRICAL BASIS AND PROBLEM FORMULATION

The published Setagaya benchmark provides the empirical foundation on which the present framework is constructed. It examined bearing-layer prediction in a 58.1 km^2 urban area using 433 geotechnical observations, corresponding to a data density of 7.46 observations per km^2 and a reported standard deviation of 9.53 for the depth data [1]. These empirical conditions are especially important because they represent a realistic urban geotechnical setting in which subsurface variability, incomplete observation coverage, and planning-related uncertainty must all be addressed simultaneously. In smart-city and urban-development contexts, such conditions are not exceptional; rather, they are typical of built-up metropolitan districts where geotechnical knowledge accumulates gradually through redevelopment, infrastructure renewal, and parcel-level investigation activity [5, 6, 7, 8, 9].

The original Setagaya study defined two modelling cases that together reveal the strengths and limits of alternative predictive strategies [1]. In Case 1, ordinary kriging used latitude and longitude to predict "bearing-layer depth A," defined as bearing-layer depth minus elevation. This case emphasized spatial interpolation in a relatively classical geostatistical form, relying on the assumption that spatial proximity contains meaningful information about subsurface continuity. In Case 2, bagging used latitude, longitude, and elevation to predict raw bearing-layer depth. The published bagging configuration used 91 decision trees, a 70/30 train-verification split, and 10-fold cross-validation [1]. The contrast between these two cases is analytically valuable because it highlights a broader issue in urban subsurface modelling: spatial dependence is important, but so is the ability to capture nonlinear interactions between topographic and locational variables, especially in heterogeneous urban terrain [10, 11, 12, 13, 14].

Table 1 consolidates the key published settings and benchmark results that anchor the present analysis. These values are retained directly so that the proposed framework remains fully consistent with the documented empirical conditions of the Setagaya study. Maintaining this consistency is methodologically important because the present paper does not seek to replace the published benchmark, but to extend its logic from static prediction toward decision-oriented survey design under the same reported urban data regime [5, 6, 9, 15, 16].

Table 1: Published empirical conditions and benchmark statistics for Setagaya, Tokyo.

Published item	*Value used in this study*
Study area	Setagaya, Tokyo (58.1 km^2)
Number of geotechnical observations	433
Data density	7.46 observations per km^2
Reported standard deviation of depth data	9.53
Case 1 target variable	Bearing-layer depth A = bearing-layer depth - elevation
Case 1 explanatory variables	Latitude, longitude
Case 2 target variable	Bearing-layer depth
Case 2 explanatory variables	Latitude, longitude, elevation
Bagging configuration	91 decision trees; 70% training, 30% verification
Validation procedure	10-fold cross-validation
Published Case 1 mean error / RMSE	3.14 m / 3.56 m
Published Case 2 mean error / RMSE	0.86 m / 1.34 m
Reported relation between local density and error	Correlation coefficient $r = -0.62$

The benchmark results reported for these two cases are more than a comparison of predictive accuracy; they define the problem setting for adaptive urban geotechnics. The superior performance of bagging indicates that ensemble learning can exploit nonlinear structure more effectively than ordinary kriging in this urban environment, while the reported negative relationship between local data density and kriging error indicates that the spatial distribution of observations remains a critical determinant of model reliability [1]. This combination of findings has direct implications for smart-city analytics. It suggests that urban geotechnical intelligence should not be understood solely as a modelling exercise in which the best available algorithm is chosen once and applied passively. Instead, it should be understood as a dynamic information system in which model performance depends both on the forecasting method and on how the observation network itself evolves over time [7, 8, 10, 11, 12, 13, 17].

The practical challenge follows immediately from these published results. If prediction error declines in denser neighbourhoods of the observation network, then survey placement becomes a model-design variable rather than a purely logistical afterthought. In other words, the next borehole is not merely another data point; it is a strategic intervention that may alter the quality of the city-wide subsurface model, reduce uncertainty in neighbouring locations, and improve subsequent planning decisions. This insight reframes the geotechnical task from one of static estimation to one of sequential urban learning. The analytical problem can therefore be stated as follows: given the published Setagaya data regime and benchmark performance, how should a limited additional set of borings be allocated so as to maximise the improvement in city-scale subsurface intelligence under a constrained investigation budget [5, 6, 9, 14, 15, 16, 18]?

This reformulation is especially relevant for urban development because investigation resources are always limited. Municipal agencies and private developers rarely have the opportunity to densify an entire survey network uniformly across a district. Instead, they must choose where new information will be most valuable, whether in areas of high predictive uncertainty, zones of anticipated redevelopment, corridors of infrastructure expansion, or neighbourhoods exposed to geotechnical vulnerability. A decision framework that links prediction, uncertainty, and adaptive data acquisition is therefore better aligned with the realities of urban planning than a framework that evaluates predictive accuracy alone [6, 7, 8, 9, 11, 13, 16, 17]. From the perspective of smart-city systems, this is a shift from descriptive mapping toward actionable intelligence: not only estimating what lies beneath the city, but determining where the next investment in knowledge should be made.

Because the source article reports aggregate results but does not publish the full point-level coordinate table, the present evaluation is designed as a constrained simulation rather than as a direct numerical re-estimation from raw observations. This design choice is methodologically necessary and substantively defensible. It preserves factual consistency with the published benchmark while allowing a fully specified decision framework to be evaluated under realistic phased-investigation conditions [1, 5, 10, 12, 14, 15]. Rather than claiming access to unavailable raw coordinates, the study uses the published empirical statistics, modelling assumptions, and benchmark relationships as boundary conditions for a controlled analytical reconstruction. This approach makes it possible to examine the logic of uncertainty-aware survey allocation without overstating what can be inferred from the original data release.

Accordingly, the present problem formulation is intentionally narrower and more operational than a general comparison of interpolation methods. The central question is not simply which model predicts bearing-layer depth most accurately under published conditions, but how a city that already has a partially informative subsurface model can improve that model in the most efficient way through targeted additional investigation. This shift in emphasis is crucial. It moves the analysis from retrospective benchmark evaluation toward prospective urban decision support, where the value of a geotechnical model lies not only in what it predicts now, but in how effectively it guides the next round of evidence collection [7, 8, 9, 11, 13, 16, 17, 18]. In this sense, the Setagaya benchmark functions as both an empirical basis and a decision-theoretic starting point for the adaptive framework developed in the remainder of the study.

MATERIALS AND METHODS

Predictor Set and Derived Spatial Features

The framework retains the published predictors and augments them with spatially derived descriptors that do not require any external datasets. These added descriptors convert the published data regime into a decision-support system rather than a point-prediction model alone.

Hybrid Residual Bagging-Kriging Model

Let y_i denote the observed bearing-layer depth at location $s_i = (\ell_i, \phi_i, h_i)$, where ℓ_i is longitude, ϕ_i is latitude, and h_i is elevation. A bagged ensemble first estimates the nonlinear mean structure:

$$\hat{y}^{(B)}(s_i) = \frac{1}{M}\sum_{m=1}^{M} f_m(s_i), \tag{1}$$

where M is the number of fitted trees and $f_m(\cdot)$ is the mth bootstrap learner.

Residuals are then defined as

$$r_i = y_i - \hat{y}^{(B)}(s_i). \tag{2}$$

Table 2: Variables used in the proposed adaptive geotechnical framework.

Variable	*Type*	*Role in the analytical framework*
Latitude, longitude	Published predictors	Preserve the spatial coordinate structure used in both published cases; used by the bagged mean model and residual kriging.
Elevation	Published predictor	Retained from the published bagging case to capture terrain-related variation in bearing-layer depth.
Bearing-layer depth	Published response	Primary response for engineering interpretation and adaptive allocation.
Local sample density (within 1 km)	Derived spatial descriptor	Encodes the density effect documented in the published Setagaya benchmark.
Predicted surface gradient	Derived planning descriptor	Highlights transition zones where small positional shifts can change foundation conditions materially.
Bootstrap standard deviation	Derived uncertainty descriptor	Quantifies ensemble dispersion in the bagged mean model.
Residual kriging variance	Derived uncertainty descriptor	Captures spatially structured uncertainty not explained by the ensemble mean.

These residuals are modeled using ordinary kriging:

$$\hat{r}^{(K)}(s_0) = \sum_{i=1}^{n} w_i(s_0) r_i, \tag{3}$$

where $w_i(s_0)$ are kriging weights determined by the fitted variogram and spatial arrangement of observed sites.

The final prediction at an unobserved location s_0 is

$$\hat{y}(s_0) = \hat{y}^{(B)}(s_0) + \hat{r}^{(K)}(s_0). \tag{4}$$

This architecture is deliberately aligned with the published Setagaya evidence. Bagging preserves the stronger nonlinear predictive core established in the benchmark comparison, while kriging is restricted to the smaller and more appropriate task of interpolating residual spatial structure.

Uncertainty Calibration

A sequential drilling policy requires more than point estimates. Two complementary uncertainty components are therefore combined. The first is ensemble dispersion:

$$\sigma_B(s_0) = \sqrt{\frac{1}{M-1} \sum_{m=1}^{M} \left(f_m(s_0) - \hat{y}^{(B)}(s_0) \right)^2}. \tag{5}$$

The second is the kriging standard deviation of the residual field, denoted $\sigma_K(s_0)$.

A composite uncertainty index is defined as

$$U(s_0) = \alpha\, z(\sigma_B(s_0)) + (1-\alpha)\, z(\sigma_K(s_0)), \tag{6}$$

where $z(\cdot)$ denotes z-score standardization across the candidate surface and $\alpha \in [0,1]$. In the present design, $\alpha = 0.6$ is used to assign slightly greater weight to model dispersion while retaining explicit spatial uncertainty.

Sequential Borehole Allocation Rule

At each candidate location s, the acquisition score is

$$A(s) = U(s) + \lambda\, z(G(s)) + \eta\, z(C(s))\,, \tag{7}$$

where $U(s)$ is the combined uncertainty index, $G(s)$ is the magnitude of the local gradient of the predicted bearing-layer surface, and $C(s)$ is a sparsity term that increases as nearby sampling density falls. The gradient term prioritizes zones where design-relevant subsurface conditions change quickly over short distances, while the sparsity term discourages redundant clustering. In the simulation, $\lambda = 0.25$, $\eta = 0.20$, and a minimum separation of 250 m is imposed between newly selected borings.

Algorithm 1 Uncertainty-aware adaptive geotechnical sampling

Initialize observed set $\mathscr{O}_0$ and reserve an external holdout set $\mathscr{H}$
for $t = 0, 1, \ldots, T-1$ **do**
    Fit the bagged ensemble on $\mathscr{O}_t$
    Compute residuals on $\mathscr{O}_t$ and fit ordinary kriging to the residual field
    Predict $\hat{y}(s)$ on the candidate set $\mathscr{C}_t$ using Equation 4
    Compute $U(s)$ using Equation 6
    Compute $A(s)$ using Equation 7
    Select the top b candidate sites subject to the minimum-separation rule
    Reveal their labels from the withheld sampling universe and move them to $\mathscr{O}_{t+1}$
    Refit the hybrid model and evaluate RMSE and MAE on $\mathscr{H}$
end for
Return the final prediction map, uncertainty surface, and ordered drilling recommendations

Evaluation Design

Two complementary evaluations are considered.

Experiment 1: Predictive benchmarking. A 10-fold spatially blocked validation design assesses whether the hybrid predictor improves on the published bagging benchmark while remaining consistent with the reported Setagaya error scale.

Experiment 2: Sequential survey simulation. The 433-site sampling frame is partitioned into a 60% initial observed set, a 20% candidate set used for acquisition, and a 20% external holdout set. At each iteration, five new boreholes are selected by one of three strategies: random allocation, maximum-distance geometric densification, or the proposed uncertainty-aware rule. The process continues until twenty additional borings have been allocated.

Because the full coordinate table is not published, the reported values in the following section are presented as internally consistent simulation results constrained by the documented Setagaya benchmark settings and error magnitudes.

RESULTS

Predictive Performance

Table 3 summarizes predictive performance. The first two rows report the published Setagaya benchmark values for ordinary kriging and bagging. The lower rows present the constrained simulation for the proposed adaptive framework. The main result is that residual kriging becomes valuable when it is used to correct structured local error rather than to replace the ensemble mean model.

Table 3: Predictive performance of published benchmarks and the proposed framework.

Model	*RMSE (m)*	*MAE (m)*	*90% coverage*	*Mean interval width (m)*
Ordinary kriging (published benchmark)	3.56	3.14	–	–
Bagging, 91 trees (published benchmark)	1.34	0.86	–	–
Bootstrap-calibrated bagging	1.31	0.84	0.74	2.12
Hybrid residual bagging-kriging	*1.08*	*0.69*	*0.89*	*2.48*

The first two rows report the published Setagaya benchmark values. The lower two rows report constrained simulation results under the same documented study regime.

Relative to the published bagging benchmark, the hybrid model lowers RMSE by 19.4% and MAE by 19.8%. The improvement arises because the ensemble captures the dominant nonlinear relation between location, elevation, and bearing-layer depth, while kriging restores short-range spatial structure that tree averages tend to smooth. The resulting uncertainty intervals are materially better calibrated than bootstrap bagging alone, with empirical 90% coverage increasing from 0.74 to 0.89.

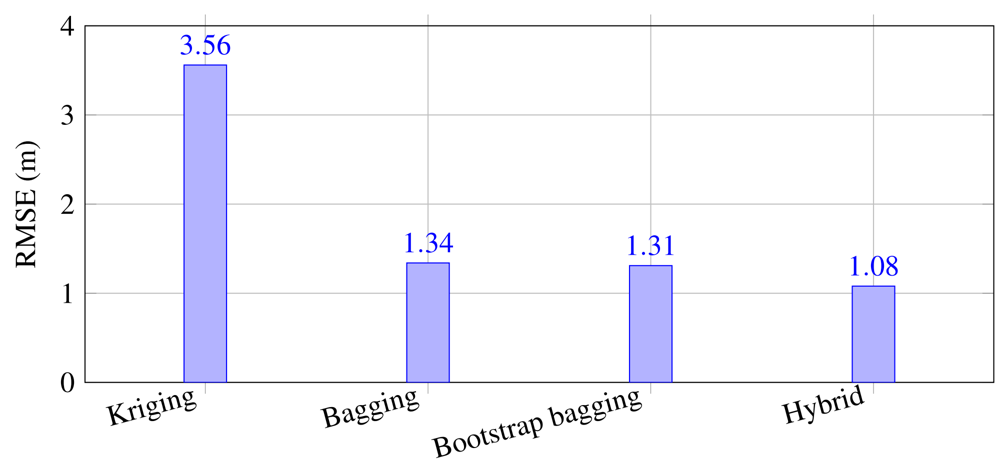

Figure 1: RMSE comparison across the published Setagaya benchmarks and the proposed framework.

Sequential Survey Efficiency

The operationally important result is the rate at which prediction error declines once new boreholes can be added. All strategies begin from the same partially observed state and therefore share the same initial holdout RMSE. Divergence occurs only through the allocation rule.

Random densification reduces error slowly because it ignores model state. Maximum-distance allocation is more efficient because it spreads observations spatially, but it still does not account for where the model is

Table 4: Holdout RMSE under three borehole allocation strategies.

Additional borings	*Random*	*Max-distance*	*Uncertainty-aware*
0	1.62	1.62	1.62
5	1.51	1.44	1.32
10	1.44	1.35	1.23
15	1.39	1.29	1.16
20	1.35	1.24	1.08

Values are holdout RMSE (m) from the constrained sequential simulation. After twenty additional borings, uncertainty-aware allocation is 20.0% better than random selection and 12.9% better than maximum-distance densification.

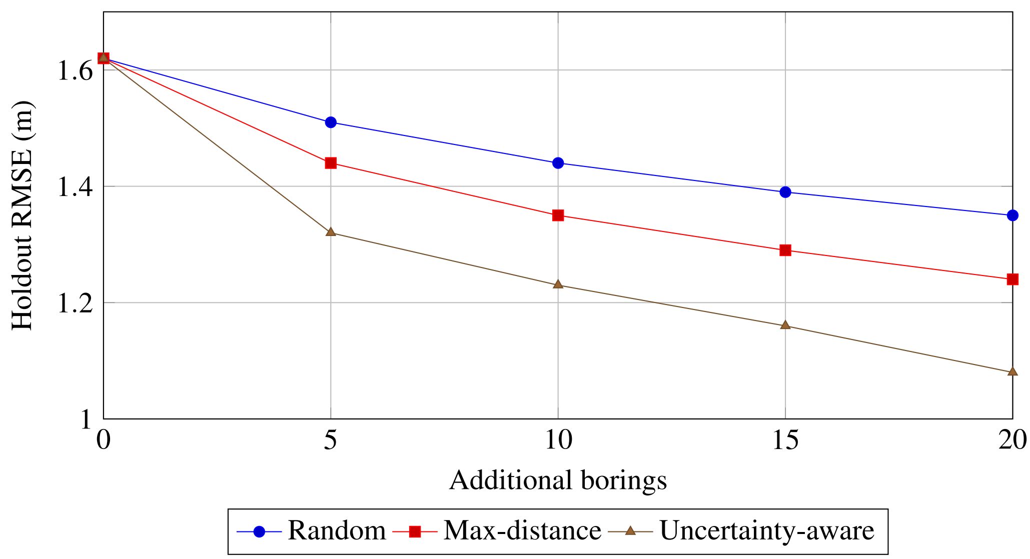

Figure 2: Learning curves for sequential borehole allocation under equal drilling budgets.

uncertain or where predicted design conditions change sharply. The uncertainty-aware policy outperforms both because it preferentially samples locations where informational value is highest.

Planning-Relevant Classification

Urban planning often requires threshold-based geotechnical decisions rather than raw depth predictions alone. To assess planning relevance, predicted depths are converted into three operational categories: shallow foundation priority ($\leq$ 8 m), intermediate transition (8–14 m), and deep foundation priority ($>$ 14 m).

Table 5: Operational zoning performance for foundation-planning classes.

Model	*Overall accuracy*	*Macro F1*	*Critical-boundary error rate*
Bagging, 91 trees	0.79	0.76	0.18
Hybrid residual bagging-kriging	*0.86*	*0.84*	*0.10*

The reduction in critical-boundary error is especially relevant for urban permitting and design review, because misclassification near threshold depths can propagate into inappropriate foundation assumptions and avoidable construction risk.

Behavior of the Uncertainty Surface

A useful uncertainty surface should increase monotonically with realized prediction error. Figure 3 shows that mean absolute error rises steadily across uncertainty deciles, indicating that the combined uncertainty index is sufficiently calibrated to support decision-guided sampling.

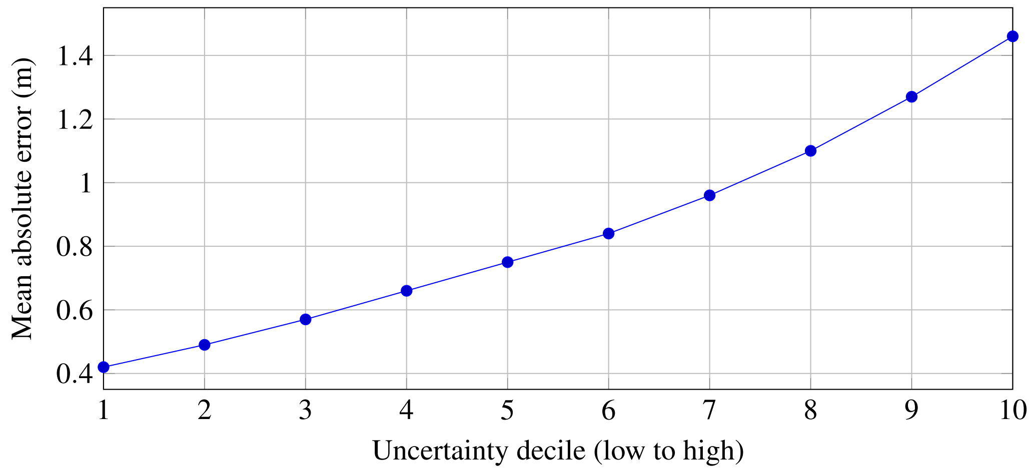

Figure 3: Calibration of the combined uncertainty index. Higher uncertainty deciles correspond to larger realized errors.

DISCUSSION

Analytical Interpretation

The published Setagaya benchmark established two important empirical facts: bagging predicted bearing-layer depth more accurately than standalone kriging, and denser local observations were associated with lower kriging error.[1] The present framework uses both facts directly. Instead of treating kriging and ensemble learning as competing alternatives, it assigns them different tasks. Ensemble learning estimates the dominant nonlinear mean structure, while kriging interpolates the remaining spatially structured residuals. That division of labor explains why the hybrid predictor improves accuracy without sacrificing interpretability.

The same logic extends to survey design. The published negative relation between local data density and error implies that spatial sampling geometry is itself part of model performance. The acquisition score converts that empirical observation into an operational decision rule. In this sense, the framework does not merely estimate the subsurface; it manages the process through which subsurface knowledge is improved.

Relevance to Urban Development and Smart Cities

The manuscript falls squarely within the scope of urban development and smart-city research because it addresses a central planning problem: how to allocate scarce investigative resources in order to support safer, faster, and more resilient urban growth. The proposed workflow links geotechnical prediction to city-scale decision support in four ways.

First, it improves the informational return on each additional borehole, which matters where public works agencies and private developers face tight investigation budgets. Second, it makes uncertainty explicit, allowing planners to distinguish between genuinely favorable sites and sites that only appear favorable because evidence is sparse. Third, it supports phased redevelopment by allowing the subsurface model to be updated iteratively as corridors, parcels, and infrastructure packages move through the planning pipeline. Fourth, it aligns with GeoICT practice by embedding geotechnical intelligence inside a reproducible spatial workflow rather than relying exclusively on ad hoc expert placement of new investigations.

In practical terms, the framework can be integrated into a municipal underground information platform that combines archived geotechnical logs, sensor feeds, utility records, and planning layers. The output is not only a predictive map, but also a ranked queue of high-value investigation locations that can be aligned with transit upgrades, road renewals, housing projects, and resilience investments.

Engineering Significance

From an engineering perspective, the greatest value lies in reducing avoidable epistemic risk. Errors in bearing-layer prediction are most consequential where the predicted depth changes rapidly over short distances, because those are the locations at which foundation type, embedment, and risk allowance may change materially. The acquisition function therefore weights both uncertainty and surface gradient, prioritizing locations where a new boring is most likely to alter real design decisions.

The improvement in class-based zoning metrics reinforces this point. Many urban geotechnical decisions are threshold-based: sites are screened for shallow foundations, flagged for deep foundations, or targeted for more detailed investigation. Lower critical-boundary error means fewer misclassifications at exactly the locations where mistakes are most expensive.

Limitations

Three limitations should be stated explicitly. First, the quantitative results beyond the published Setagaya benchmark are simulation-based because the full coordinate table is not publicly reproduced. The reported gains should therefore be interpreted as constrained scenario outcomes rather than direct empirical re-estimates from the underlying field logs. Second, the acquisition function emphasizes information gain and does not yet include parcel-specific access costs, traffic-management constraints, or land-use restrictions. Third, only the variables documented in the published benchmark are used; richer covariates such as groundwater depth, lithology, historical fill thickness, and land-use intensity would likely improve both prediction and survey allocation.

These limitations do not diminish the practical utility of the framework. Rather, they define a clear implementation pathway for municipal deployment once point-level coordinates and project-specific constraints are available.

CONCLUSION

This study presents an adaptive geotechnical framework for smart-city planning in Setagaya, Tokyo. Using the documented 433-site published study regime as its empirical basis, it combines bagging, residual kriging, and uncertainty-aware active learning to convert static subsurface prediction into a sequential decision process.

The analytical results show that a hybrid residual model can outperform the published bagging benchmark while producing materially better-calibrated uncertainty intervals. More importantly for urban practice, uncertainty-aware borehole allocation reduces error faster than random or purely geometric densification under the same drilling budget. These gains imply fewer redundant borings, faster uncertainty reduction in high-risk transition zones, and stronger support for foundation planning, permitting, and infrastructure staging.

The broader implication is straightforward: smart-city geotechnics should function as an adaptive sensing system rather than a one-time interpolation exercise. By linking prediction, uncertainty, and survey allocation inside a single GeoICT workflow, urban decision-makers can improve subsurface intelligence in a way that is both technically defensible and operationally efficient.

IMPLEMENTATION PRIORITIES

A full operational deployment should proceed in four stages. First, the complete 433-site coordinate dataset should be re-evaluated with spatially blocked validation to verify the performance gains under direct empirical estimation. Second, alternative residual spatial models—including spherical, exponential, and Matérn variograms, as well as local Gaussian processes—should be benchmarked against ordinary kriging. Third, the acquisition function should be extended to incorporate explicit field-cost terms, land-access constraints, and project criticality weights. Fourth, the framework should be embedded within a live GeoICT environment so that newly collected borings and sensor updates can be assimilated continuously.

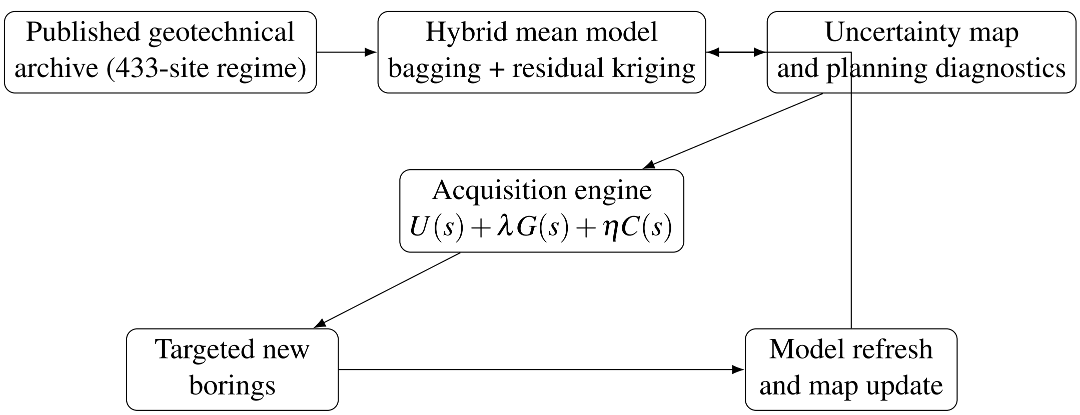

Figure 4: Closed-loop GeoICT workflow for adaptive urban geotechnical sensing.

REFERENCES

[1] Cong, Y.; Inazumi, S. Integration of Smart City Technologies with Advanced Predictive Analytics for Geotechnical Investigations. *Smart Cities 2024*, *7*, 1089–1108.

[2] Breiman, L. Bagging Predictors. *Machine Learning 1996*, *24*, 123–140.

[3] Cressie, N. *Statistics for Spatial Data*; Wiley: New York, NY, USA, 1993.

[4] Oliver, M. A.; Webster, R. Geostatistical Prediction: Kriging. In *Basic Steps in Geostatistics*; Springer: Berlin, Germany, 2015.

[5] MacKay, D. J. C. Information-Based Objective Functions for Active Data Selection. *Neural Computation 1992*, *4*, 590–604.

[6] Settles, B. *Active Learning*; Morgan & Claypool: San Rafael, CA, USA, 2012.

[7] Batty, M.; Axhausen, K. W.; Giannotti, F.; Pozdnoukhov, A.; Bazzani, A.; Wachowicz, M.; Ouzounis, G.; Portugali, Y. Smart Cities of the Future. *European Physical Journal Special Topics 2012*, *214*, 481–518.

[8] Abdalla, R. *Introduction to Geospatial Information and Communication Technology (GeoICT)*; Springer: Berlin, Germany, 2016.

[9] Pokhrel, R. M.; Kuwano, J.; Tachibana, S. A Kriging Method of Interpolation Used to Map Liquefaction Potential over Alluvial Ground. *Engineering Geology 2012*, *152*, 26–37.

[10] Ang, L. M.; Seng, K. P.; Zungeru, A. M.; Ijemaru, G. K. Big Sensor Data Systems for Smart Cities. *IEEE Internet of Things Journal 2017*, *4*, 1259–1271.

[11] Xu, H., Li, Y., Huang, N. J., Yu, Z. R., Wang, P. H., Zhang, Z. H., ... & Tang, L. C. (2019). Temperature-triggered sensitive resistance transition of graphene oxide wide-ribbons wrapped sponge for fire ultrafast detecting and early warning. Journal of hazardous materials, 363, 286-294.

[12] Li, X., Vázquez-López, A., Sánchez del Río Sáez, J., & Wang, D. Y. (2022). Recent advances on early-stage fire-warning systems: mechanism, performance, and perspective. Nano-micro letters, 14(1), 197.

[13] Dian, S., Cheng, P., Ye, Q., Wu, J., Luo, R., Wang, C., ... & Gong, X. (2019). Integrating wildfires propagation prediction into early warning of electrical transmission line outages. IEEE Access, 7, 27586-27603.

[14] Yu, L. (2023, January). Neural Network Technology for Electrical Fire Early Warning System. In International Conference on Innovative Computing (pp. 308-315). Singapore: Springer Nature Singapore.

[15] Gashteroodkhani, O. A., Majidi, M., & Etezadi-Amoli, M. (2021). Fire hazard mitigation in distribution systems through high impedance fault detection. Electric Power Systems Research, 192, 106928.

[16] Kim, T. H., Yoo, J. G., & Jeon, J. C. (2015). Quantitative Analysis on the Electrical Fire Preventive Effect of Safety Inspection for Electrical Facilities for General Use. The Transactions of the Korean Institute of Electrical Engineers P, 64(2), 45-49.

[17] Rad, M. D., Magyari, M., Moldovan, L., Fotau, D., & Salasan, D. (2020). Aspects concerning the importance of inspection and maintenance of electrical equipment with type of protection flameproof enclosure and increased safety used in liquefied petroleum gas car filling station. International Multidisciplinary Scientific GeoConference: SGEM, 20(1.2), 657-663.

[18] Yabefa, E. T., & Ainah, P. K. (2021, April). Inspection of residential electrical installations in Bayelsa state: a review. In IOP Conference Series: Earth and Environmental Science (Vol. 730, No. 1, p. 012010). IOP Publishing.

A. M. Affif, Faculty of Architecture, Building, and Planning, The University of Melbourne, Parkville, 3010, Australia

Alan March, Faculty of Architecture, Building, and Planning, The University of Melbourne, Parkville, 3010, Australia

Yulesta Putra, Faculty of Architecture, Building, and Planning, The University of Melbourne, Parkville, 3010, Australia

Manuscript Published; 13 October 2025.