Street Art & Urban Creativity | Vol. 11, No. 4, 2025 | ISSN 2183-9956
Street Art & Urban Creativity. Scientific Journal

A Nationwide Predictive Digital-Twin Framework for Resilient Transportation, Bridges, and Public Infrastructure

CHAITANYA NAGABHYRU[1]
[1]12-1006, RAHMATH NAGAR 2ND LINE, CHILAKALURIPET, GUNTUR PIN: 522616, ANDHRA PRADESH, INDIA
CHAITANYANAGABHYRU@GMAIL.COM

Abstract

Public agencies must decide each year which bridges, culverts and pavements to inspect and which to repair, using inspection records that are infrequent, sensors that cover only a few assets, and data from other regions that may not resemble their own. This paper presents a predictive digital-twin framework that (i) keeps a continuously updated estimate of each asset's health and its uncertainty from inspections, sensors, traffic and weather, (ii) predicts which assets will reach a critical deficiency within two years, (iii) transfers what is learned in data-rich regions to regions with few labeled assets, and (iv) reports uncertainty that grows honestly when data are missing, so that uncertain assets are routed to inspection. Because agency records cannot be shared, the framework is evaluated on a transparent synthetic fleet of 2200 bridges in six regions, four with abundant history and two, with older and more corrosion-exposed stock, treated as data-scarce targets, repeated over 5 simulations. On unseen target assets the framework reaches an area under the ROC curve of 0.941 with 100 labeled target assets, against 0.811 for a standard condition-based model, and cuts calibration error from 0.078 to 0.045 relative to a source-only model. Inspecting the riskiest 10% of assets captures 52% of critical events, against 22% when ranking by the latest inspection rating. In a five-year budget simulation the framework lowers total cost by 48% relative to reactive repair and 32% relative to condition-based maintenance. Most of the gain comes from the digital-twin health estimate, whereas transfer mainly improves calibration, and all results are simulation-based.

Keywords: Digital Twin, Predictive Maintenance, Bridge Management, Transfer Learning, Uncertainty Quantification, Inspection Prioritization, Infrastructure Resilience.

1. Introduction

The United States relies on more than six hundred thousand bridges in the National Bridge Inventory, together with millions of kilometers of pavement, culverts, tunnels and transit assets. Federal standards require routine bridge inspections at intervals of up to twenty-four months, and each inspection yields condition ratings on a scale from zero to nine for the deck, superstructure and substructure [1]. These records are the backbone of asset management, yet they are sparse in time, subject to inspector judgment, and silent about what happens between visits. Recent federal investment in bridges and roads has raised expectations that agencies will spend public money where it prevents the most failures, which requires knowing where risk is concentrated.

Three obstacles stand in the way. First, data are incomplete. Inspection records may be missing or late, only a small share of structures carry sensors, and sensors drop out. A prediction that ignores missing data is overconfident, and a prediction that treats it as zero is biased. Second, data are unevenly distributed across places. A state that has inspected thousands of bridges for decades holds far more labeled history than a neighboring agency with a few hundred, and the two may differ in climate, traffic, materials and age, so a model fitted on one may fail on the other. Third, decisions are consequential and made under budget limits. A ranking of assets by risk is useful only if the risk is calibrated and if the agency knows which rankings are uncertain.

A digital twin offers a natural organizing idea: a living model of each asset that is updated

whenever new evidence arrives and that carries its own uncertainty forward [7], [8]. This paper develops and evaluates a framework built around that idea. The objectives are: (i) to construct, for every asset, a digital state estimate of health and deterioration from inspections, sensors, traffic and environmental data, with uncertainty that widens when data are absent; (ii) to predict critical deficiency within two years; (iii) to transfer predictive knowledge from data-rich regions to data-scarce ones and to measure how much local labeled data are needed; (iv) to report uncertainty that supports decisions, including deferral of uncertain assets to inspection; and (v) to quantify the consequence for inspection prioritization and maintenance cost. Performance is measured by failure prediction, maintenance-cost reduction, inspection prioritization and model transferability.

The topic is of national interest: longer-lasting infrastructure, safer transportation and more efficient public spending. Digital twins and predictive maintenance are named among current federal science priorities, and a framework that is honest about its uncertainty is one that agencies can adopt without overstating what the data support. Because realistic labeled records from agencies are not public, all experiments use a synthetic benchmark whose construction is described in Section V, and results are properties of the method on that benchmark. The paper is organized as follows: Section II reviews the literature, Section III positions related work, Section IV presents the methodology, Section V the experimental setup, Section VI the results, Section VII the discussion and limitations, and Section VIII the conclusions.

2. LITERATURE REVIEW

Bridge deterioration has long been modeled as a stochastic process. Reviews of life-cycle performance models describe Markov chains, gamma processes and reliability-based approaches that predict condition states and their uncertainty over time [2]. Markov-chain models of bridge decks estimate transition probabilities between condition ratings from inspection histories and are widely used in management systems [3]. Their limitation is that transition probabilities are usually estimated per population and do not adapt to an individual asset's exposure, sensor data or inspection history. Data-driven alternatives such as neural networks learn deterioration from many covariates [4], and dynamic Bayesian belief networks combine expert knowledge and data while representing uncertainty about condition [5].

Structural health monitoring adds measurements of vibration, strain and displacement to condition ratings, from which damage-sensitive features can be derived [6]. In practice, monitoring is limited to a minority of structures because of cost, so predictive systems must work with a mixture of instrumented and uninstrumented assets. The digital-twin concept, introduced in product lifecycle management, keeps a virtual representation synchronized with the physical object [7], and surveys of industrial twins emphasize real-time synchronization, fidelity and integration with decisions [8]. A classic tool for synchronization with noisy, irregular measurements is the Kalman filter, which tracks a hidden state and its uncertainty and automatically widens that uncertainty when measurements are missing [9].

Transfer between assets or regions is a problem of dataset shift. Surveys of transfer learning distinguish settings by whether labels are available in the target domain and by what changes between domains [10]. When only the distribution of inputs changes, reweighting source examples by the density ratio between target and source corrects the bias [11], although weights become unstable when the two distributions barely overlap [12]. When labeled target data exist but are scarce, common approaches fine-tune or recalibrate a source model, pool the data with extra weight on the target, or train only on the target. Which works best depends on the size of the target sample and the similarity of the domains.

Uncertainty estimation for predictions is addressed by ensembles, whose disagreement approximates model uncertainty [13], by approximate Bayesian methods [14], and by conformal prediction, which converts any predictor into one with distribution-free coverage guarantees under exchangeability [15], [16]. Gradient-boosted trees remain strong tabular predictors and handle missing values natively [17]. Probability calibration, the agreement between predicted and observed frequencies, is distinct from discrimination and matters when probabilities feed cost calculations [18], and the Brier score combines both aspects [21]. Missing data raise the further issue of whether the missingness mechanism itself is informative [19]. Finally, selective prediction studies when a classifier should abstain, trading coverage for accuracy [20].

3. RELATED WORK

The closest prior work comes from three directions. Condition-prediction studies in bridge management fit deterioration models to inspection data, with the best-known approaches based on Markov chains and neural networks [3], [4]. They usually treat each region or agency separately, assume that the inspection record is complete, and rarely report uncertainty in a form that can drive decisions. Dynamic Bayesian approaches provide uncertainty but typically require expert-specified structures and have been demonstrated on limited samples [5].

Digital-twin studies for civil structures tend to focus on high-fidelity physical models of individual assets, such as finite-element models updated with sensor data, which is costly to build and does not scale to a fleet in which most assets are uninstrumented [6], [8]. At the other end, data-only machine-learning studies scale to many assets but discard the dynamic, uncertainty-bearing structure that a twin provides. The framework proposed here combines a lightweight state-estimation twin for every asset with a learned risk model, so that the learning step receives uncertainty-aware health estimates as inputs.

TABLE 1. Positioning of the proposed framework relative to common practice

Capability	Common practice	Treatment in this framework
Health estimate	Latest inspection rating	Filtered health state with uncertainty, using inspections, sensors and work orders
Incomplete data	Imputation or exclusion	Uncertainty widens when data are missing; missingness handled natively
Cross-region learning	Separate model per agency	Source-region model, reweighting and recalibration on few target labels
Uncertainty report	Point estimate or none	Ensemble spread, twin variance and scaled conformal intervals
Planning	Fixed cycles or worst-first	Risk-ranked maintenance and uncertainty-first extra inspections

The contribution is therefore the integration of a fleet-scale twin, a transfer step and uncertainty reporting under one evaluation, with attention to the data-scarce target, to missing data and to downstream cost. No component is claimed to be new in isolation. As in related work, the field lacks open, labeled fleet data with failure outcomes, which limits what any single study can demonstrate.

4. METHODOLOGY

Framework Overview

Figure 1 shows the five modules. Records of inspections, sensors, traffic, weather and work orders feed a digital twin that estimates each asset's health and deterioration rate with uncertainty (Module 1). A failure model trained on data-rich source regions predicts the probability that an asset reaches a critical deficiency within two years (Module 2). A transfer step adapts that model to a target region using whatever labeled assets exist there (Module 3). An uncertainty module combines the spread among ensemble members, the twin's own uncertainty and conformal intervals to report how much each prediction can be trusted (Module 4). A planning module uses risk and uncertainty to allocate a limited number of preventive maintenance actions and extra inspections (Module 5). New inspections and work orders flow back to the twin each year.

.

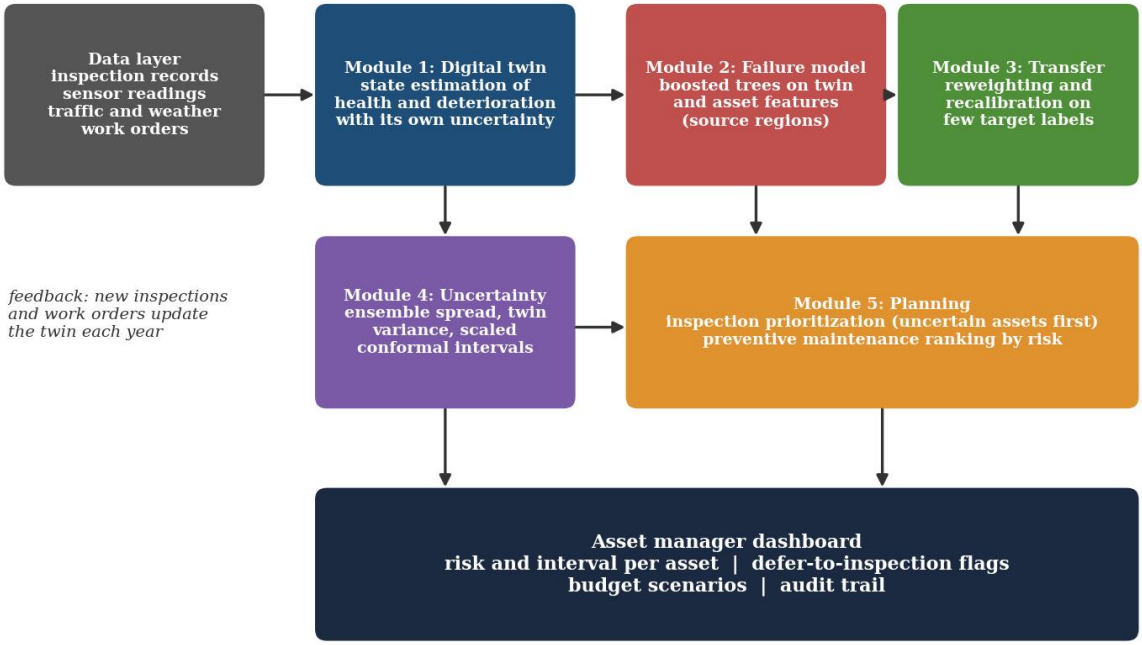


Figure 1: Architecture of the proposed predictive digital-twin, transfer and uncertainty framework.

Digital Twin of Each Asset

Each asset's twin tracks two hidden quantities: a health index from zero to one hundred and its yearly rate of change. The twin is a linear state-space filter in the Kalman family [9]. Every year it first moves its estimate forward, assuming health continues to change at the current rate and allowing a small random drift in both level and rate, which increases its uncertainty. It then corrects the estimate with whatever measurements arrived that year. An inspection rating, multiplied by eleven to place it on the health scale, is treated as a noisy reading of health, and a sensor index is treated as a second, less noisy reading when available. When no measurement arrives, the correction step is skipped and uncertainty simply grows, which is how the twin represents incomplete data. A recorded repair resets the estimate to a high value with small uncertainty, and a preventive maintenance action raises the estimate and slows the assumed rate. The prior for a new asset depends on its age. The three noise settings were chosen from a small grid by the accuracy of one-year-ahead rating predictions on source-region assets only.

From the filtered state the twin produces a forecast of health two years ahead with its standard deviation, the implied probability that health will fall below the critical threshold of forty, and a predictive distribution for the next inspection rating, which can be checked against reality when the inspection occurs. These outputs, together with the current health estimate, its uncertainty and the estimated rate, become inputs for the failure model.

Failure Prediction

The prediction task is to estimate, at the end of each year, the probability that an asset reaches a critical deficiency in either of the next two years. Inputs combine static attributes (age, material class, traffic, chloride exposure), recent environment (average freeze-thaw exposure over three years, flood events), inspection history (latest rating, years since inspection, change between the last two ratings), sensor summaries, time since maintenance, and the twin outputs. A gradient-boosted tree classifier [17] is used because it handles missing values natively and captures interactions. Region identity is deliberately excluded so that the model can be applied to a region it has never seen.

Transfer to a Data-Scarce Region

The transfer step has three parts. First, source examples are reweighted by how typical they are of the target region, using a regularized classifier that distinguishes source from target assets on age, traffic, climate and material, with weights limited to a narrow range to avoid instability

[11]. Second, instead of one model, an ensemble of twelve boosted classifiers is trained on bootstrap resamples of the reweighted source data; the average of their predictions in log-odds form is the base risk, and their disagreement is the model-uncertainty signal [13]. Third, when labeled target assets are available, a small, strongly regularized logistic model is fitted on the target labels, taking the ensemble's base risk and three covariates (age, chloride exposure and freeze-thaw exposure) as inputs. This corrects shifts in the overall event rate and in the effect of exposure while leaving the learned structure intact. With no target labels the ensemble's uncertainty-aware base risk is used unchanged.

Four alternatives serve as baselines: a source-only model without adaptation; a model trained only on the labeled target assets; a pooled model trained on source and target assets together with each target row counted five times; and a standard condition-based logistic regression that uses only age, material and the latest rating. Two further baselines isolate the role of the twin: a model that uses the twin's raw forecast probability with no learning, and a boosted model trained without any twin outputs.

Uncertainty Reporting

Three kinds of uncertainty are reported. The twin's own standard deviation measures how well the asset's health is known and grows when data are missing. The spread of the ensemble measures how much plausible models disagree on the risk. For the next inspection rating, three forecast intervals are compared: the twin's raw ninety percent interval, a fixed-width conformal interval whose half-width is the ninetieth percentile of absolute errors on calibration assets, and a scaled conformal interval whose half-width is the calibrated ninetieth percentile of errors divided by the twin's predicted standard deviation, multiplied by the standard deviation for the asset in question [15], [16]. The scaled version is designed so that intervals widen automatically when the twin is less certain. Calibration uses the observed inspection ratings of the 400 pool assets in the training years, with no data removed; it needs ratings but not failure outcomes.

For decisions, a combined uncertainty score adds the ensemble spread and the twin's standard deviation, each divided by its median. Assets with the highest scores can be deferred to physical inspection instead of relying on the prediction, a form of selective prediction [20].

Planning Policies

The planning module is evaluated by simulation. In each year, with a fixed budget, it chooses assets for preventive maintenance and assets for an extra inspection. The proposed policy maintains the assets with the highest adapted risk and spends its extra inspections on the assets with the highest combined uncertainty among those not selected. A risk-averse variant ranks by an upper plausible risk, adding one ensemble standard deviation. Comparison policies are no preventive work (reactive), maintenance of the oldest assets, maintenance of the assets with the lowest latest rating, maintenance by the source-only model, and an unattainable oracle that knows true health. All policies other than the proposed ones allocate extra inspections to the assets longest without inspection. Assets maintained within the previous four years are excluded from selection.

Evaluation Measures

Failure prediction is measured by the area under the ROC curve, the average precision, the Brier score, and an equal-frequency calibration error computed over ten bins of predicted probability. Inspection prioritization is measured by the share of critical events captured when only a given fraction of assets, ranked each year, is inspected first. Transferability is measured by how metrics change with the number of labeled target assets and by the gap between source-region and target-region performance. Cost is measured over a five-year horizon.

5. EXPERIMENTAL SETUP

Synthetic Fleet

The benchmark simulates 2,200 bridges in six regions over fifteen years (2010–2024). Each asset has a latent health index that falls each year at a rate that depends on its material (steel, reinforced concrete or prestressed concrete), on chloride exposure, on freeze-thaw exposure, on

truck traffic and, increasingly, on its age; the rate varies randomly from year to year, and occasional flood shocks cause sudden drops. Before 2010, assets are aged forward from construction with agency repairs when health becomes poor, so that health at the start of the record is realistic. An asset reaches a critical deficiency when its health falls below forty; it is then repaired to a high level, and the event is recorded. Preventive maintenance raises health by twelve points and slows deterioration by a quarter for four years.

Observations are imperfect. Every asset is inspected every two years, with half the fleet inspected in even years and half in odd years; twelve percent of inspection records are missing; and each rating is the true health divided by eleven, plus an asset-specific bias and random noise, rounded to a whole number between zero and nine. Thirty percent of source-region assets and only fifteen percent of target-region assets have a sensor that reports health with its own bias and noise; sensor readings drop out in twenty percent of years, and sensors fail permanently at random times. Traffic and weather are observed.

Regions differ systematically (Table 2). Four source regions each contain 400 assets. The two target regions each contain 300 assets and differ from the source regions in ways that shift both the inputs and the relationship between inputs and outcomes: the coastal region has older stock, very high chloride exposure and a stronger effect of chloride on deterioration, and the northern region has high freeze-thaw exposure with a stronger effect of freezing. The share of assets reaching a critical deficiency within two years is therefore about 3.9% in source-region training windows and 13.4% in target test windows.

TABLE 2. Regions in the synthetic fleet

Region	Role	Mean age (years)	Chloride exposure (0–1)	Freeze-thaw index	Relative deterioration
R0 temperate	Source	38	0.25	40	1.00
R1 cold, de-icing	Source	42	0.35	80	1.05
R2 warm	Source	30	0.22	5	0.90
R3 mountain	Source	36	0.30	100	1.10
R4 coastal, older stock	Target	52	0.75	30	1.20 (stronger chloride effect)
R5 northern	Target	47	0.50	120	1.15 (stronger freeze effect)

.

Data Splits and Training Protocol

For each simulation, the 600 target assets are randomly divided into 200 test assets and a pool of 400 from which labeled target assets are drawn in nested subsets of 0, 25, 50, 100, 200 and 400. Training windows are years 2011–2016, whose outcomes are known by 2018. Test windows are years 2018–2022, whose outcomes extend to 2024, giving 1,000 test windows per simulation with about 134 critical events. Unlabeled target covariates, with no outcomes, are used for reweighting. Five independent simulations were run, each with a new fleet, new noise, a new target split and new model fits, and results are reported as mean ± standard deviation across simulations. Boosted models use 150 iterations, a learning rate of 0.06, at most 15 leaves per tree, a minimum of 40 examples per leaf and a mild regularization penalty; implementation uses scikit-learn [22].

TABLE 3. Principal experimental settings

Component	Setting
Fleet	2,200 assets; 4 source regions (400 each); 2 target regions (300 each)
Period and prediction task	2010–2024; critical deficiency (health below 40) within next two years
Observation process	Biennial inspections, 12% records missing; sensors on 30% (source) and 15% (target) of assets
Twin settings	Level noise 0.50, rate noise 0.01, inspection noise 8 health points (grid-selected on source assets)
Transfer sweep	0, 25, 50, 100, 200, 400 labeled target assets; 5 simulations
Ensemble	12 bootstrap boosted classifiers on reweighted source windows
Missing-data stress	25%, 50%, 75% of inspection and sensor observations removed at test time
Planning budget	Per year, for 200 assets: 10 preventive maintenance actions and 20 extra inspections
Costs	Preventive action $0.20M; critical event $2.4M (emergency repair and user cost); extra inspection $4,000

Budget Simulation

Policies are compared on the 200 test assets from 2020 to 2024. All start from the same

simulated state at the end of 2019, and the random inputs (deterioration noise, floods, inspection noise) are identical across policies, so differences arise only from decisions. Models used by the policies are trained on information through 2018 and the twin is updated each year from the observations the policy generates, including those from extra inspections. Maintenance and inspection allocations are made at the start of each year from information available at the end of the previous year.

6. RESULTS

Failure Prediction and the Role of the Digital Twin

Table 4 reports performance on unseen target assets with 100 labeled target assets available. A standard condition-based logistic model, using age, material and the latest rating, reaches an area under the ROC curve of 0.811 ± 0.016. Adding the full feature set to a boosted model without any twin outputs raises it only to 0.833 ± 0.011. Including the twin's health estimate, rate and forecast lifts the source-only boosted model to 0.940 ± 0.010. The twin's raw two-year forecast probability, with no learning at all, scores 0.945 ± 0.006, which is as good as any learned model. The twin is thus responsible for most of the gain, because it combines noisy inspections and sensors over time and knows how uncertain it is, whereas the latest rating is a single noisy reading. The average precision shows the same ordering (Table 4).

TABLE 4. Failure prediction on unseen target assets (100 labeled target assets where applicable; mean ± sd over 5 simulations)

Method	ROC AUC	Average precision	Brier score	Calibration error
Condition-based logistic	0.811 ± 0.016	0.306 ± 0.025	0.110 ± 0.004	0.074 ± 0.007
Twin forecast only (no learning)	0.945 ± 0.006	0.693 ± 0.033	0.098 ± 0.005	0.099 ± 0.006
Source-only boosted model, no twin inputs	0.833 ± 0.011	0.338 ± 0.028	0.104 ± 0.004	0.062 ± 0.006
Source-only boosted model	0.940 ± 0.010	0.692 ± 0.045	0.085 ± 0.007	0.078 ± 0.014
Source-only, importance-weighted	0.939 ± 0.011	0.704 ± 0.039	0.086 ± 0.010	0.079 ± 0.020
Target-only boosted model	0.925 ± 0.013	0.624 ± 0.061	0.079 ± 0.007	0.051 ± 0.009
Pooled (target rows counted five times)	0.942 ± 0.009	0.679 ± 0.042	0.075 ± 0.009	0.051 ± 0.013
Proposed (transfer + uncertainty)	0.941 ± 0.008	0.701 ± 0.048	0.068 ± 0.006	0.045 ± 0.013

Removing the twin inputs from the full proposed pipeline reduces the area under the curve from 0.941 to 0.833 ± 0.016. Figure 2 illustrates a twin for one bridge: its estimate follows the hidden true health through noisy inspections and sensor readings, and its uncertainty band widens in years with no data and narrows when a measurement arrives. The figure also shows how the target regions differ from the source regions in age and chloride exposure.

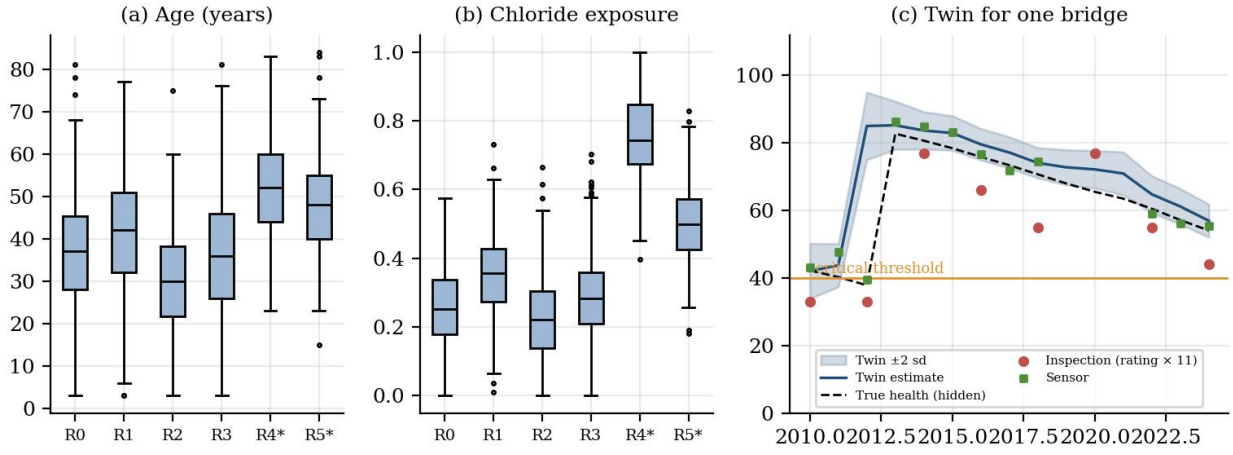


Figure 2: (a, b) distribution of age and chloride exposure by region (* target regions); (c) twin estimate with two-standard-deviation band for one bridge

Model Transferability

The source-only boosted model reaches an area under the curve of 0.946 on held-out windows in its own source regions and 0.940 on the target regions, a gap of only 0.006. Discrimination therefore transfers almost without loss in this benchmark, because the twin summarizes health in a region-independent way. What does not transfer is the event rate: the target regions have a much higher share of critical events, and the source-only model underestimates risk. Its calibration error is 0.078, and Figure 4 shows that its predicted

probabilities fall well below observed frequencies in the upper bins. Reweighting the source examples to resemble the target did not help (area under the curve 0.939, calibration error 0.079); the target regions overlap little with the source regions, and the reweighting retained only about a fifth of the effective sample size in the first simulation.

Figure 3 and Table 5 show how methods respond to labeled target data. With no target labels, the proposed model matches the source-only model on discrimination (0.943). A target-only model needs many labeled assets: with 25 labeled assets it reaches 0.904, with 100 it reaches 0.925, and only with 200 to 400 does it match the transferred models (0.937 and 0.943). The pooled model is the strongest baseline in discrimination at small sample sizes, with 0.939 at 25 labeled assets compared with 0.932 for the proposed model. The proposed model has the best Brier score at every sample size from 25 upward (for example 0.070 at 25 labeled assets against 0.079 for the pooled model), and its calibration error falls from 0.089 with no labels to 0.045 with 25, a level the pooled model reaches only with about 200 labeled assets. The differences in area under the curve among adapted models are smaller than the variation across simulations, so the honest summary is that adaptation mainly improves probability accuracy and calibration, not ranking.

TABLE 5. Effect of the number of labeled target assets (mean over 5 simulations; AUC / calibration error)

Labeled target assets	Target-only	Pooled	Proposed
0	–	–	0.943 / 0.089
25	0.904 / 0.043	0.939 / 0.064	0.932 / 0.045
50	0.912 / 0.048	0.942 / 0.056	0.937 / 0.046
100	0.925 / 0.051	0.942 / 0.051	0.941 / 0.045
200	0.937 / 0.044	0.941 / 0.045	0.941 / 0.042
400	0.943 / 0.040	0.946 / 0.042	0.943 / 0.043

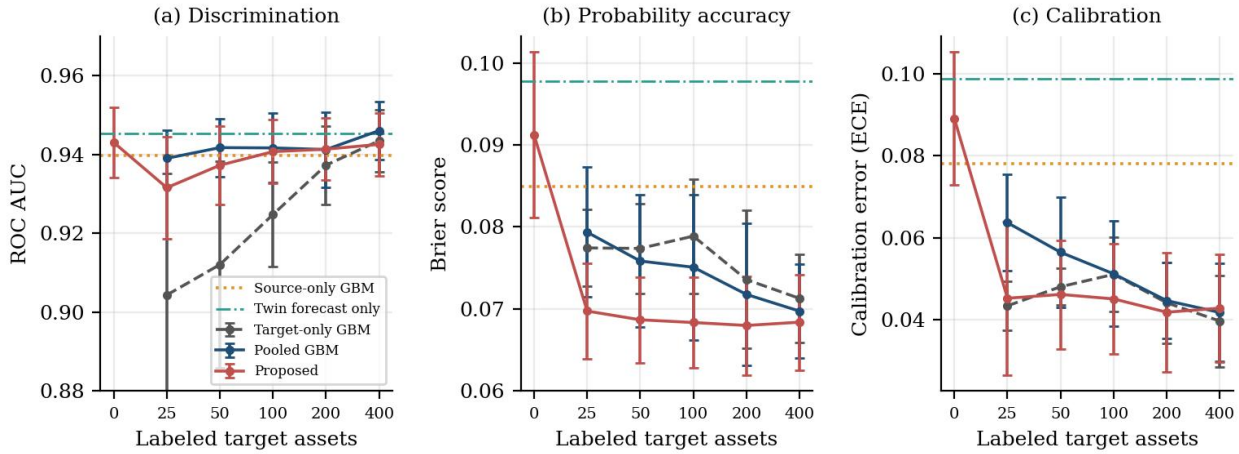


Figure 3: Effect of the number of labeled target assets on (a) discrimination, (b) Brier score and (c) calibration error

Performance by target region is similar. In the coastal region, with 14.1% of windows ending in a critical event, the area under the curve is 0.937 for the source-only model and 0.941 for the proposed model. In the northern region, with 12.6%, the values are 0.945 and 0.945.

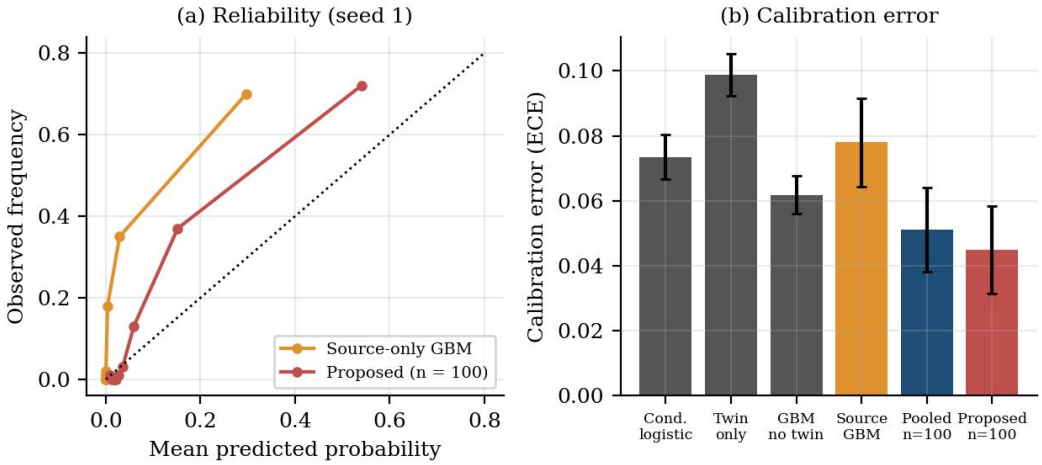


Figure 4: (a) reliability of predicted probabilities for the source-only and proposed models in one simulation; (b) calibration error by method

Reporting Uncertainty under Incomplete Data

Table 6 reports what happens when increasing shares of inspection and sensor observations are removed at test time, while models remain trained on complete data. Discrimination degrades gracefully, with the area under the curve falling from 0.941 with no removal to 0.915 at fifty percent and 0.901 at seventy-five percent removal, and calibration error stays between 0.045 and 0.058. The uncertainty reports respond as intended: the twin's average standard deviation rises from 4.3 to 6.7 health points, and the ensemble spread rises from 0.52 to 0.58.

TABLE 6. Uncertainty reporting as observations are removed (target test assets; mean over 5 simulations)

Observations removed	AUC	Calibration error	Twin sd (health points)	Ensemble sd (log-odds)	Coverage: raw twin (%)	Coverage: fixed conformal (%)	Coverage: scaled conformal (%)
0%	0.941	0.045	4.3	0.52	73.3	85.0	83.0
25%	0.934	0.047	4.7	0.54	72.3	84.0	82.8
50%	0.915	0.053	5.3	0.55	72.3	81.3	82.8
75%	0.901	0.058	6.7	0.58	71.3	78.4	81.7

For the next inspection rating, the twin's raw ninety percent interval covers only about 73% of realized ratings, even with complete data, so the filter is overconfident, partly because ratings are rounded integers with asset-specific biases that the filter does not model. Conformal calibration improves coverage to 85% (fixed width) and 83% (scaled), still below the nominal ninety percent, because calibration used earlier years and test years differ (exchangeability does not hold exactly). The more important difference is robustness to missing data. As observations are removed, the coverage of the fixed-width interval falls to 78% and that of the raw twin interval stays near 71%, while the scaled interval holds at 82% by widening from 3.48 to 4.07 rating points. No interval attained its nominal coverage on the test years, and Figure 5 shows these results alongside the width and selective prediction analyses.

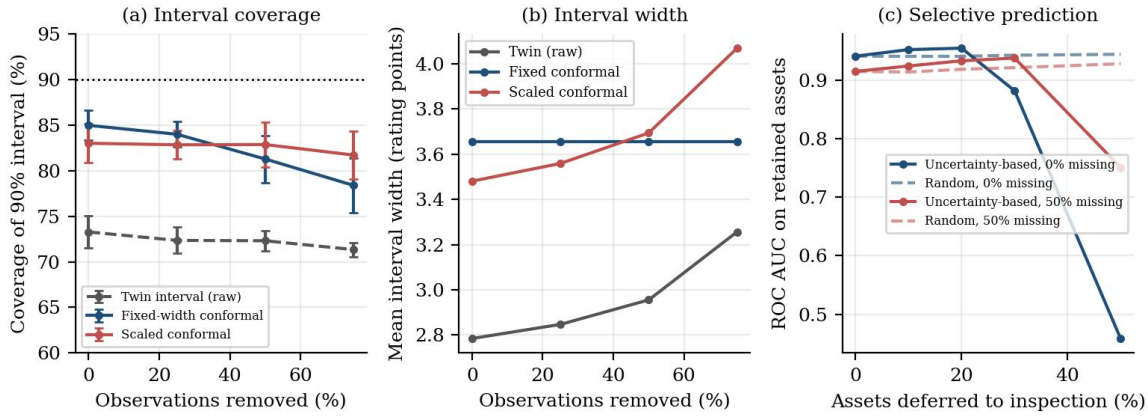


Figure 5: (a) coverage and (b) width of 90% intervals for the next inspection rating as observations are removed; (c) discrimination on retained assets when the most uncertain are deferred to inspection

Table 7 and Figure 5(c) examine selective prediction. Deferring the ten and twenty percent of assets with the highest combined uncertainty to physical inspection raises the area under the curve on the remaining assets from 0.941 to 0.952 and 0.954, whereas random deferral leaves it unchanged. With half of the observations removed, deferring thirty percent lifts it from 0.915 to 0.937. Deferral has a limit: with complete data, deferring thirty percent or more removes many of the highest-risk assets, because uncertainty and risk are correlated, and the area under the curve on the retained assets then falls (0.459 at fifty percent deferral, with large variation across simulations). Deferral should therefore be used for the most uncertain tenth to fifth of assets, and the deferred assets should be inspected, not ignored.

TABLE 7. Selective prediction: AUC on retained assets after deferring the most uncertain (or random) share

Share deferred	0% missing: uncertainty-based	0% missing: random	50% missing: uncertainty-based	50% missing: random
0%	0.941	0.941	0.915	0.915
10%	0.952	0.940	0.924	0.914
20%	0.954	0.940	0.933	0.918
30%	0.882	0.942	0.937	0.921
50%	0.459	0.944	0.751	0.928

Inspection Prioritization

Figure 6 and Table 8 show how well each method concentrates critical events among the assets inspected first. Random selection captures about 10% of events when inspecting ten percent of assets, ranking by age about 12%, and ranking by the latest inspection rating about 22%. The twin forecast alone captures 50%, the source-only boosted model 50% and the proposed model 52%. At twenty percent, the corresponding values are 48% for the latest rating and 82% for the proposed model. The three twin-informed methods are indistinguishable in prioritization, so the added value of transfer here is calibration, not ranking.

TABLE 8. Share of critical events captured when the top-ranked assets are inspected first (mean ± sd over 5 simulations)

Ranking method	Top 5% (%)	Top 10% (%)	Top 20% (%)	Top 30% (%)
Random	4.8 ± 1.8	9.7 ± 1.4	19.8 ± 2.8	30.0 ± 3.1
Age	6.8 ± 1.4	12.1 ± 1.9	23.0 ± 3.9	30.0 ± 4.9
Lowest latest rating	11.4 ± 1.7	21.7 ± 1.7	48.1 ± 3.1	70.0 ± 3.1
Twin forecast only	27.7 ± 1.2	50.3 ± 2.4	82.6 ± 3.4	96.2 ± 1.4
Source-only GBM	28.9 ± 1.4	50.4 ± 3.7	81.1 ± 4.4	95.8 ± 2.1
Proposed (transfer + uncertainty)	28.8 ± 2.0	51.5 ± 3.0	82.1 ± 2.3	94.8 ± 1.6

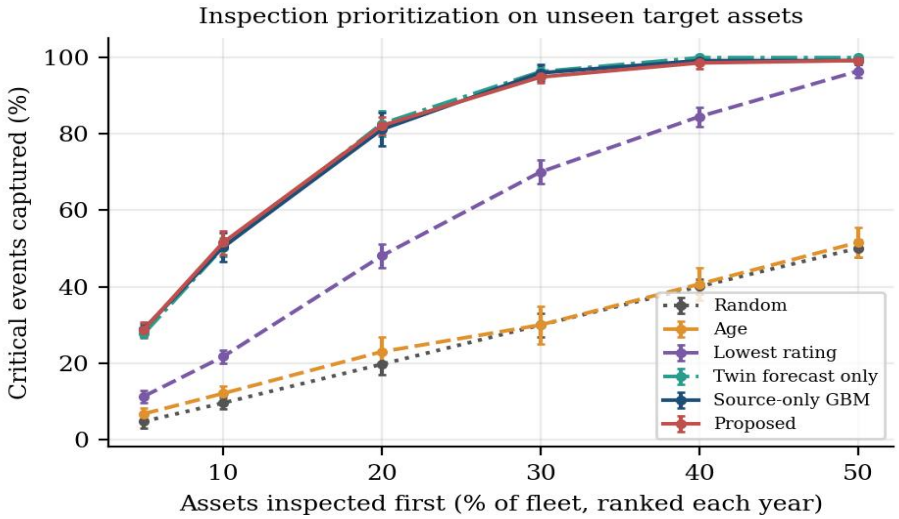


Figure 6: Share of critical events captured when the top-ranked share of assets is inspected first

Maintenance Cost

Table 9 and Figure 7 give five-year results for the 200 test assets under the common budget of ten maintenance actions and twenty extra inspections per year. With no preventive work, the fleet experiences 66.8 critical events and costs $160.3 million. Maintaining the oldest assets

reduces cost by only 7.5% because age is a weak predictor in this fleet. Maintaining assets with the lowest latest rating reduces cost by 23.6%. Risk-based maintenance with the source-only model reduces cost by 46.0%, and the proposed policy by 48.0%, which is 31.7% below the condition-based policy and avoids about 36 critical events. An oracle with knowledge of true health reaches 64.0%, so the proposed policy realizes roughly three quarters of the attainable reduction.

The difference between the proposed policy and the source-only policy is small. The proposed policy is cheaper by $3.4 million on average, but the paired difference has a standard deviation of $6.2 million across the five simulations and cannot be distinguished from zero. The risk-averse variant, which ranks by an upper plausible risk, is equivalent. In other words, in this benchmark the benefit of using the twin and risk model is large and clear, while the additional benefit of transfer adaptation and uncertainty-guided inspection is not demonstrated on cost.

TABLE 9. Five-year budget simulation for 200 target assets (mean ± sd over 5 simulations)

Policy	Critical events	Total cost ($M)	Cost reduction vs reactive (%)	Cost reduction vs condition-based (%)	True mean health, end of 2024
Reactive (no preventive work)	66.8 ± 3.9	160.3 ± 9.3	0.0 ± 0.0	-31.4 ± 8.7	61.8 ± 0.5
Time-based (oldest first)	57.4 ± 3.1	148.2 ± 7.4	7.5 ± 1.9	-21.5 ± 7.3	63.7 ± 0.4
Condition-based (lowest rating)	46.6 ± 2.9	122.2 ± 6.9	23.6 ± 4.9	0.0 ± 0.0	61.5 ± 0.7
Source-only ML	31.8 ± 4.1	86.7 ± 9.8	46.0 ± 4.3	29.0 ± 7.2	58.9 ± 0.5
Proposed (transfer + uncertainty)	30.4 ± 4.0	83.4 ± 9.7	48.0 ± 5.2	31.7 ± 7.8	58.8 ± 0.9
Proposed, risk-averse	30.2 ± 4.3	82.9 ± 10.3	48.4 ± 4.8	32.2 ± 8.2	58.6 ± 0.8
Oracle (true health)	19.8 ± 3.7	57.9 ± 8.9	64.0 ± 3.7	52.5 ± 7.7	56.2 ± 0.7

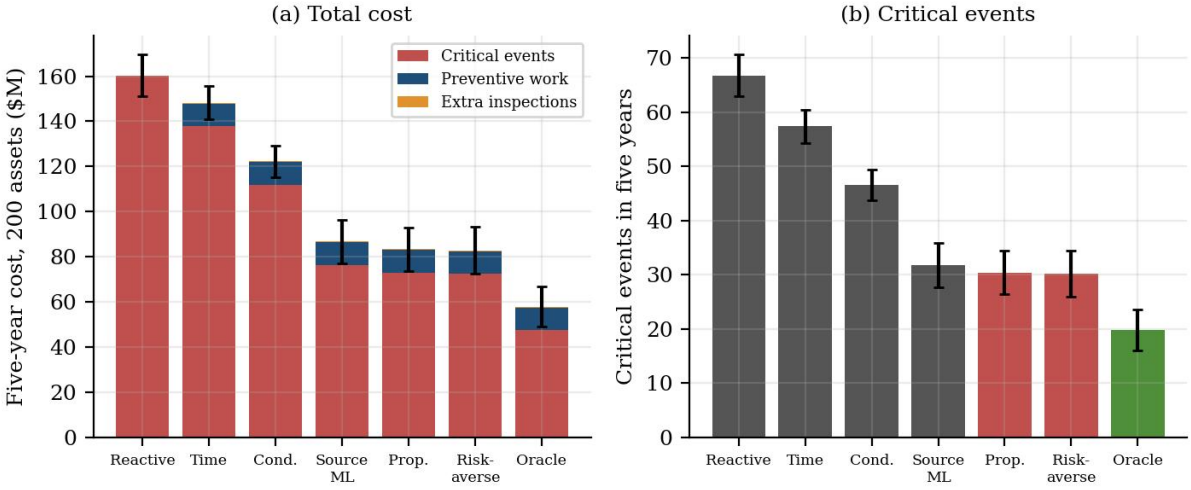


Figure 7: (a) five-year cost by component and (b) critical events for each maintenance policy

7. DISCUSSION

Four conclusions emerge. First, the greatest value comes from the digital twin: a filtered health estimate that uses every observation and knows its own uncertainty. A boosted model without it performed only a little better than a condition-based logistic model, and a twin forecast with no learning matched the best learned model. This is plausible on engineering grounds, since an inspection rating is a noisy snapshot and the twin integrates evidence over time. It is also a reason for caution, because the twin's linear assumptions match the way the simulator generates data, and real deterioration is more complex and may need richer state models.

Second, transfer is easiest where the physics are shared. Discrimination transferred with a gap of only about six thousandths in area under the curve, and the failure of source-only models was in calibration, because target regions have more frequent events. A modest number of local labels, around twenty-five to fifty assets in this benchmark, was enough to repair calibration and probability accuracy, whereas a model trained on local data alone needed about two hundred assets to match discrimination. This points to a practical recommendation for agencies with limited records: borrow a model from data-rich neighbors, then use whatever local outcomes exist to recalibrate it, instead of either using the borrowed model as is or building one from

scratch. The result should not be read as a general guarantee. The regions in the benchmark were designed to share one deterioration mechanism, and a target with a different mechanism, for example a different structure type, would be a harder test.

Third, uncertainty is useful when it is connected to action. The twin's variance and the scaled conformal interval respond to missing data, while a fixed-width interval does not, and deferring the most uncertain tenth to fifth of assets to inspection improved discrimination on the rest. None of the intervals reached nominal coverage on the test years, which shows the practical difficulty of calibration when the data distribution changes over time, and the uncertainty-guided deferral hurts when pushed too far because uncertainty and risk are correlated. In practice, intervals should be recalibrated each year and monitored.

Fourth, the cost results are large relative to simple rules, with a reduction of about a third compared with condition-based maintenance, but small relative to the oracle gap and indistinguishable between risk-based policies with and without the transfer step. Spending money on better risk estimates is clearly worthwhile, whereas the marginal value of calibrated, uncertainty-aware ranking depends on how the costs are structured. If decisions were threshold-based, for example approving a repair when risk exceeds a cost-based cutoff, calibration would matter more than it does for fixed-budget ranking. The framework supports both uses, and the choice should follow the agency's decision process. These findings are consistent with the emphasis in current federal priorities on digital twins and predictive maintenance, in that the gains arise from integrating the data an agency already collects, with no need for universal sensing.

The limitations are substantial. The fleet is synthetic and was generated by processes the authors specified, so absolute values for accuracy and cost are not predictions for real agencies. Only bridges are modeled; transfer between asset classes, such as from bridges to culverts or pavements, was not tested. Failure is defined as crossing a health threshold, whereas real failures involve specific mechanisms, scour, fatigue and overload, that the model does not distinguish. The inspection ratings do not reflect inspector disagreement beyond a fixed bias and noise, and missing data were removed at random, not systematically, as in real agencies where the missingness may depend on asset condition. The cost model uses assumed unit costs and does not include traffic delay costs beyond a lump sum. Only five simulations were run, and the differences among adapted models were smaller than their variation. Future work should validate on real inspection histories from partner agencies, test transfer across asset types and across structures with different mechanisms, model informative missingness, and study how asset managers respond to uncertainty reports.

8. CONCLUSION

This paper presented a predictive digital-twin framework for infrastructure that estimates asset health with uncertainty, predicts critical deficiency, transfers knowledge to data-scarce regions and reports uncertainty that responds to missing data. On a synthetic fleet of 2,200 bridges in six regions, the framework reached an area under the ROC curve of 0.941 on unseen target assets with 100 labeled target assets, compared with 0.811 for a condition-based model. It captured 52% of critical events among the top tenth of assets, against 22% for ranking by the latest rating, and lowered five-year cost by 48% compared with reactive repair and 32% compared with condition-based maintenance. The digital twin accounted for most of this benefit. Transfer adaptation improved calibration and probability accuracy, especially with fewer than one hundred labeled assets, but not discrimination or cost, and scaled conformal intervals stayed closer to nominal coverage than the alternatives as data became sparse, without reaching it.

These conclusions describe a controlled benchmark. The most useful next steps are validation on real inspection records, tests of transfer across asset types, richer deterioration models inside the twin, and joint studies with asset managers of how uncertainty reports change decisions, so that the promise of predictive maintenance for safer and more durable public infrastructure is tested against operational evidence.

References

1. A Mahendra Vardhan and S. Sridhar, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree

Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458

2. Dip Bharatbhai Pate, Dynamic vs. Static visualizations: Understanding their use cases in big data analysis, International Journal of Science and Research Archive, 2022, 05(02), 392-395., DOI https://doi.org/10.30574/ijsra.2022.5.2.0077

3. Dip Bharatbhai Patel, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021,DOI: https://doi.org/10.30574/ijsra.2021.2.1.0047
4. Dip Bharatbhai Patel, Understanding the Impact of Text Analytics on Social Media Sentiment Analysis, Vol. 5 No. 5 (2019) ,September-October, DOI: https://doi.org/10.32628/CSEIT1936413
5. Dip Bharatbhai Patel, Real-Time Data Streaming in BI Dashboards: Opportunities and Challenges, Archives / Vol. 5 No. 4 (2019) : July-August, DOI: https://doi.org/10.32628/CSEIT192142
6. Dip Bharatbhai Patel, the Role of Data Visualization in Modern Intelligence, Archives / Vol. 2 No. 2 / Research Article, Published, 2020-11-28 DOI: https://doi.org/10.32996/jcsts.2020.2.2.4
7. A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: https://ijcnis.org/index.php/ijcnis/article/view/8682
8. A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4347/3189
9. A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4603
10. Backup Integrity and Recovery Readiness Assessment for High-Availability Databases Raghu Gollapudi, Computer Fraud & Security, Vol. 2023, Issue 8, pp. 35-48, 10.52710/cfs.1031
11. Risk-Controlled Near-Zero-Downtime Oracle Database Migration Using GoldenGate Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 8(3), 2022 ,10.22399/ijcesen.5382
12. Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep),Publication year 2021, Volume 2021, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/970
13. Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS),Publication year 2023, Volume 2023, Issue 12,https://computerfraudsecurity.com/index.php/journal/article/view/971
14. Manoj Kumar Kagitha, LOAD BALANCING TECHNIQUES IN DISTRIBUTED CLOUD ENVIRONMENTS, (IJESMR - 2017), D.O.I:10.53555/jcr.v6:i1.13380, https://jcreview.com/archives/volume-6/issue-1/13380
15. Manoj Kumar Kagitha, AI-DRIVEN OBSERVABILITY FOR HYPERSCALE COLOCATION DATA CENTERS: A CASE STUDY OF LOG ANALYTICS AND ANOMALY DETECTION PIPELINES, Archives / Vol. 27 No. 7 (2019): JOCAAA,https://www.eudoxuspress.com/index.php/pub/article/view/5053
16. Manoj Kumar Kagitha, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, https://www.eudoxuspress.com/index.php/pub/article/view/5054
17. Naresh Lokiny. (2022). Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. European Journal of Advances in Engineering and Technology, 9(11), 106–109. https://doi.org/10.5281/zenodo.13325989
18. Naresh Lokiny, (2021), "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", International Journal of Science and Research (IJSR), 10(3), 2024-2027. https://dx.doi.org/10.21275/SR24724151733, https://www.ijsr.net/getabstract.php?paperid=SR24724151733
19. Naresh Lokiny, & Ranganath Nandanampati. (2020). DevSecOps: Integrating Security into

DevOps with AI in Cloud. Journal of Scientific and Engineering Research, 7(10), 239–242. https://doi.org/10.5281/zenodo.13348695
20. Naresh Lokiny, & Pradip Reddy. (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. European Journal of Advances in Engineering and Technology, 8(3), 69–72. https://doi.org/10.5281/zenodo.13325845
21. Jayanth Para, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/171 , DOI: https://doi.org/10.46121/pspc.51.4.3
22. Jobayar Alom, Ahsan Ahmed. (2023). Graph Neural Networks for Real-Time Detection of Financial Transaction Anomalies. Acta Scientiae, 24(5), 82–90. https://doi.org/10.22178/acta.24.5.6
23. Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
24. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
25. Vikas Suresh Bagora, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/359
26. Vikas Suresh Bagora, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/362
27. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611 , https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611
28. Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods, 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP)
https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ & citation_for_view=efbNMkwAAAAJ:M3NEmzRMIkIC
29. Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426
30. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611, , https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611
31. Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
32. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
33. Kaleshwar Aryasomayajula, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415,

https://pspac.info/index.php/dlbh/article/view/374
34. Controlled Synthesis and Characterization of Lanthanum Nanorods, Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan, doi:10.18576/ijtfst/090205, URL: https://www.naturalspublishing.com/Article.asp?ArtcID=20705, May-2020
35. Sudipkumar Ghanvat , Aishwarya Badlani, Shreya Sandeep Joshi, Marketing Effectiveness in the Post-Cookie Era: A Comparative Analysis of First Party and Zero-Party Data Strategies on Campaign Performance and Customer Equity, Vol. 31 No. 4 (2023): JOCAAA , https://www.eudoxuspress.com/index.php/pub/article/view/3940
36. Chander Vijay S Sanbhi, BAYESIAN UPDATING OF RAM MODELS FOR DYNAMIC AVAILABILITY FORECASTING UNDER REALISTIC DOWNTIME CONSTRAINTS, Vol. 51 No. 3 (2023): July-September 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/381
37. Chander Vijay S Sanbhi, HYBRID RELIABILITY MODELLING FOR ROTATING EQUIPMENT: PHYSICS- INFORMED LEARNING WITH SPARSE FAILURE DATA, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/382
38. A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: https://ijcnis.org/index.php/ijcnis/article/view/8682
39. A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4347/3189
40. A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4603
41. Santhosh Chitraju Gopal Varma, The Evolution of Cloud-Native Architectures: Exploring the Synergy between Kubernetes and Microservices, DOI: https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P104, Vol. 1 No. 4 (2020), https://ijetcsit.org/index.php/ijetcsit/article/view/108
42. Santhosh Chitraju Gopal Varma, Artificial Intelligence in Cloud Computing: Building Intelligent, Distributed, and Fault-Tolerant Systems, Vol. 3 No. 1 (2022), DOI: https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P105, https://ijaibdcms.org/index.php/ijaibdcms/article/view/78
43. Sunil Karmakar, AI driven Identity & Access Management with Oracle Cloud security, Vol. 30 No. 2 (2022): JOCAAA, https://eudoxuspress.com/index.php/pub/article/view/4118
44. Bandari, V. V., & Lekkala, H. B. (2024). Physics-informed reinforcement learning for real-time control of complex manufacturing processes. Power System Protection and Control, 52(2), 164–170. https://pspac.info/index.php/dlbh/article/view/343
45. Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. Power System Protection and Control, 52(1), 38–45. https://pspac.info/index.php/dlbh/article/view/342
46. Kancharla Vijayalakshmi, Bethapudi Samuel Anand Andrews, Estimation of Cyproheptadine Hydrochloride and Tricholine Citrate Simultaneously in Syrup Dose by Applying Stability Indicating HPLC Methodology, December 31, 2020, Vol. 11, Issue 4, Pages 1813–1819, DOI: 10.26452/ijrps.v11iSPL4.4383, https://ijrps.com/v11si4/3683
47. Vijayalakshmi Kancharla ,Kumar Abbaraju V D N, Andrews B S A, Nagaraju Dasari, Nagu Korupolu, Melphan as Anticancer and New Insights of Drug as Parenteral Dosage form against SARS-CoV-2 by Using Development and Validation , November 18, 2020, Vol. 11, Issue 1, Pages 1437–1446
DOI: 10.26452/ijrps.v11iSPL1.3686, https://ijrps.com/v11si1/2911
48. Ramesh, D., Singh, A., & Channaiah, L. H. (2024). The importance of validated kill steps and thermal death kinetics of Salmonella in food safety management. Global Journal of Nutrition & Food Science, 5(1), Article 000601. https://doi.org/10.33552/GJNFS.2024.05.000601
49. Bharati Reddi, Chandan Kishor,Aruna Jangam, Sandeepchowdary Bala, Uma Rajeswari Batchu ,Rambabu Gundla , Anthony Addlagatta, Regioselectivity in inhibition of peptide deformylase from Haemophilus influenzae by 4- vs 5-azaindole hydroxamic acid derivatives: Biochemical, structural and antimicrobial studies, Date: November 2022, Article: 106095, Volume: Volume 128, 10.1016/j.bioorg.2022.106095,

https://www.sciencedirect.com/science/article/abs/pii/S0045206822005016
50. Bhavita Kattula, Bharati Reddi, Aruna Jangam, Lekhika Naik , Bala Manikanta Adimoolam , Suresh Vavilapalli , Sayanna Are , Jagadeshwar Reddy Thota, Surender Singh Jadav , Mohammed Arifuddin Anthony Addlagatta, Development of 2-chloroquinoline based heterocyclic frameworks as dual inhibitors of SARS-CoV-2 MPro and PLPro, Date: 1 July 2023 Article: 124772, Volume: Volume 242, Part 2, 10.1016/j.ijbiomac.2023.124772 https://www.sciencedirect.com/science/article/pii/S0141813023016665
51. Iram Irfan, Asghar Ali, Bharati Reddi, Mohd. Abrar Khan, Phool Hasan, Sarfraz Ahmed Amad Uddin, Magdalena Piatek, Kevin Kavanagh, Mohammad Abid, Design, Synthesis and Mechanistic Studies of Novel Isatin-Pyrazole Hydrazone Conjugates as Selective and Potent Bacterial MetAP Inhibitors, 19 August 2022, https://doi.org/10.3390/antibiotics11081126 https://www.mdpi.com/2079-6382/11/8/1126
52. GAURAV KUMAR SINHA, Efficient Test-Time Compute Scaling for Resource-Constrained Large Language Models Vol. 32 No. 1 (2024): JOCAAA, DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458 https://eudoxuspress.com/index.php/pub/article/view/5918
53. GAURAV KUMAR SINHA , The Impact of Artificial Intelligence on Predictive Quality Assurance in Modern Manufacturing, Vol. 33 No. 05 (2024): JOCAAA, https://doi.org/10.1109/ICAC3N56670.2022.10074458 , https://eudoxuspress.com/index.php/pub/article/view/5919
54. GAURAV KUMAR SINHA , Hybrid Quantum-Classical Neural Networks for Combinatorial Optimization and Machine Learning, Vol. 33 No. 08 (2024): JOCAAA, DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458, https://eudoxuspress.com/index.php/pub/article/view/5921
55. Lalith Chandra Bandaru , Mohammed Shakeer Bandrevu, GRAPHSEC: Graph-Based Supply Chain Attack Detection and Risk Propagation Analysis for Enterprise Salesforce Deployment Pipelines, Paper Id: 233135,Published On: 2023-11-03, Published In: Volume 11, Issue 6, November-December 2023, https://www.ijirmps.org/research-paper.php?id=233135 , DOI: https://doi.org/10.37082/IJIRMPS.v11.i6.233135