

分类号 P237

密级 公开

UDC 550

编号 10533

中南大学

博士后研究工作报告

城市地面沉降的超前预测：从机器学习到物理引导的人工智能

Kouadio Kouao Laurent

工作起止日期 2022年7月—2026年5月

报告提交日期 2026年3月

中南大学（湖南）

2026年3月

城市地面沉降的超前预测：从机器学习到物理引导的 人工智能

Advanced Forecasting of Urban Land Subsidence: From Machine Learning to Physics-Guided AI

博士后姓名	Kouadio Kouao Laurent
流动站（一级学科）名称	地质资源与地质工程
专 业（二级学科）名称	地球探测与信息技术
流动站所在学院	地球科学与信息物理学院
合作导师	柳建新教授
工作起止日期	2022 年 7 月—2026 年 5 月

中南大学人事处（湖南）

2026 年 3 月

工作报告原创性声明

本人郑重声明，所呈交的工作报告是本人在合作导师指导下进行的研究工作及取得的研究成果。尽我所知，除了工作报告中特别加以标注和致谢的地方外，报告中不包含其他人已经发表或撰写过的研究成果，也不包含为获得中南大学或其他教育机构的学位或证书而使用过的材料。与本人共同工作的同志对本研究所作的贡献，均已在文中作了明确说明。

工作报告与资料若有不实之处，本人承担一切相关责任。

作者签名：_____ 日期：_____年____月____日

工作报告版权使用授权书

本工作报告作者和合作导师完全了解中南大学有关保留、使用工作报告的规定：即学校有权保留并向国家有关部门或机构送交工作报告的复印件和电子版；本人允许本工作报告被查阅；学校可以将本工作报告的全部或部分内容编入有关数据库进行检索，可以采用复印、缩印或其他手段保存和汇编本工作报告。

保密工作报告待解密后适用本授权书。

作者签名：_____ 导师签名：_____

日期：_____年____月____日 日期：_____年____月____日

城市地面沉降的超前预测：从机器学习到物理引导的人工智能

摘要: 地面沉降是快速城市化与地下水高强度开发区域典型地质灾害，会直接破坏基础设施、降低城市排水能力，并加剧沿海地区内涝与相对海平面上升风险，威胁城市安全与可持续发展。受地下水动态、可压缩地层厚度、城市建设、降雨补给及区域水文地质条件耦合控制，地面沉降具有强非线性、空间异质性与时间滞后性，当前仍缺乏兼具高精度、可解释性与不确定性量化能力的预测模型。

本研究围绕城市地面沉降预测构建“数据驱动模拟—物理引导预测”两阶段研究框架，系统探索机器学习与物理约束融合的建模路径。第一阶段以广州市南沙区为研究区，采用极端梯度提升回归 (XGBR)、长短期记忆网络 (LSTM)，融合多源地学与城市环境变量开展沉降模拟。结果显示，地下水埋深与建筑密度是控制区域沉降演化的核心因素，机器学习模型可较好还原地面沉降的时空分布特征，证明异质城市与地质数据能够为沉降预测提供有效支撑。

纯数据驱动模型存在物理机制解释不足、长期预测易偏离地质规律、不确定性难以合理表达等局限。为此，第二阶段提出面向多步预测的物理引导人工智能框架 GeoPriorSubsNet，将概率型时空预测网络与简化水文-地质力学约束深度融合，在预测累计沉降量的同时，同步反演有效水力传导系数、有效比储系数及松弛时间尺度，使模型兼具统计精度、物理一致性与机制可解释性。

该框架在水文地质条件差异显著的李沙、中山两大城市盆地完成系统验证。验证结果表明，GeoPriorSubsNet 预测性能优异：南沙区决定系数 $R^2=0.88$ ，平均绝对误差 9.27 mm；中山区 $R^2=0.90$ ，平均绝对误差 1.78 mm。相比无物理约束模型，该框架在预测精度、概率区间覆盖率、区间锐度及物理一致性上均更具优势。其中，物理约束对结构清晰的中山地区精度与不确定性质量提升显著，对异质性更强的南沙地区则主要增强模型稳健性与可解释性。研究揭示，南沙与中山分属两类典型城市沉降响应体系：南沙驱动变量复杂、空间异质性强、不确定性区间宽；中山水文地质结构清晰、驱动-响应关系紧致。模

型反演的物理参数场与区域地质规律一致，表明物理约束的作用具有明显区域依赖性。外部验证显示，模型推断的有效压缩厚度与钻孔资料一致性良好，有效渗透性与抽水试验差异则反映井点尺度与网格尺度的固有区别。

跨城市迁移实验表明，模型直接零样本迁移效果有限，经热启动（warm start）微调后可快速恢复目标区域预测性能，证明 GeoPrior-SubsNet 是具备物理结构的可迁移先验框架，而非固定通用模型，为多城市推广应用提供可行路径。

本研究实现了城市地面沉降预测从纯数据驱动机器学习向物理引导人工智能的跨越，所提框架兼顾预测精度、不确定性表达与物理可解释性，可为地面沉降风险识别、长期演化评估、地下水管控、基础设施规划及城市韧性建设提供量化技术支撑，具有重要理论价值与实践意义。

图 10, 表 1, 参考文献 71

关键词：地面沉降；沉降预测；机器学习；物理引导人工智能；InSAR；地下水开采；水文地质建模；不确定性量化；城市风险

分类号： P237； TP181

Advanced Forecasting of Urban Land Subsidence: From Machine Learning to Physics-Guided AI

Abstract: Urban land subsidence is a major geohazard in rapidly urbanizing and groundwater-stressed regions because it threatens infrastructure stability, drainage systems, and long-term flood resilience. Forecasting this process remains difficult because subsidence is governed by coupled hydrogeological, geomechanical, climatic, and anthropogenic controls that are spatially heterogeneous, temporally lagged, and strongly nonlinear. This study develops a two-stage framework for urban land subsidence forecasting. The first stage investigates whether machine-learning models can recover subsidence dynamics from heterogeneous urban and environmental observations. Using Nansha as a case study, models based on eXtreme Gradient Boosting Regressor (XGBR) and Long Short-Term Memory (LSTM) networks show that groundwater depth and building concentration are dominant controls and that data-driven methods can provide accurate subsidence simulation. The second stage addresses the main limitation of purely data-driven prediction by introducing a physics-guided artificial-intelligence framework, GeoPriorSubsNet, which combines probabilistic spatio-temporal learning with reduced hydro-geomechanical constraints. The model jointly forecasts cumulative subsidence while inferring effective hydraulic conductivity, specific storage, and a closure-consistent relaxation timescale. The framework was evaluated in two contrasting urban basins, Nansha and Zhongshan. It achieved strong predictive skill in both cities, with $R^2 = 0.88$ and a mean absolute error of 9.27 mm in Nansha, and $R^2 = 0.90$ with a mean absolute error of 1.78 mm in Zhongshan. The results further show that the physics-guided framework improves the balance between deterministic accuracy, probabilistic calibration, and physical consistency relative to non-physics baselines. Cross-city experiments indicate that direct transfer is limited under basin shift, whereas warm-start adaptation recovers most of the target-city skill. These findings show that urban land subsidence forecasting can be advanced

from accurate machine-learning simulation toward a more interpretable, uncertainty-aware, and physically constrained modeling framework, with direct relevance for urban risk assessment, groundwater management, and infrastructure planning.

There are 10 figures, 1 tables, and 71 references in this report.

Keywords: land subsidence; subsidence forecasting; machine learning; physics-guided AI; InSAR; groundwater depletion; hydrogeological modeling; uncertainty quantification; urban risk

符 号 表

AI	Artificial intelligence
InSAR	Interferometric synthetic aperture radar
LSTM	Long short-term memory
XGBR	eXtreme Gradient Boosting Regressor
TFT	Temporal Fusion Transformer
HALNet	Hybrid attentive recurrent learning network
PINN	Physics-informed neural network
GWL	Groundwater level
$s(i, t)$	Cumulative land subsidence at grid cell i and time t
$d(i, t)$	Groundwater depth, taken as positive downward
$h(i, t)$	Hydraulic head reconstructed from surface elevation and groundwater depth
$h_{\text{ref}}(i)$	Reference hydraulic head used to define drawdown
$\Delta h(i, t)$	Drawdown relative to the reference hydraulic head
$R(i, t)$	Rainfall forcing
$U^*(i, t)$	Urban-load or urbanisation forcing proxy
$z_{\text{surf}}(i)$	Surface elevation at grid cell i
$H_{\text{eff}}(i)$	Effective compressible thickness
$H_d(i)$	Effective drainage thickness
$c_H(i)$	Censor indicator associated with the thickness field
$x_i^{(\text{sta})}$	Static input vector
$X_{i,t}^{(\text{dyn})}$	Dynamic historical input tensor

$X_{i,t}^{(\text{fut})}$	Future-known input tensor
$C_{i,t}$	Forecast-aligned coordinate and context features
Ω	Spatial domain or set of grid cells
L	Look-back window length
H	Forecast horizon length
\mathcal{Q}	Set of forecast quantile levels
q	Individual quantile level
$Q(i, t_h)$	Effective source or sink term in the reduced groundwater equation
$K(i, t_h)$	Effective hydraulic conductivity
$S_s(i, t_h)$	Effective specific storage
$\tau(i, t_h)$	Learned relaxation timescale
$\tau_{\text{cl}}(i, t_h)$	Closure-implied consolidation timescale
$\delta \log \tau(i, t_h)$	Learned log-timescale discrepancy relative to the closure relation
α_H	Drainage-thickness ratio linking H_d to H_{eff}
κ_b	Dimensionless closure-consistency factor
$R_{\text{gw}}(i, t_h)$	Reduced groundwater-flow residual
$R_{\text{cons}}(i, t_h)$	Reduced consolidation residual
$R_{\text{prior}}(i, t_h)$	Prior-mismatch residual in log-timescale space
L_{data}	Total probabilistic data loss
$L_{\text{gw}}, L_{\text{cons}}, L_{\text{prior}}$	Groundwater residual loss, consolidation residual loss, and timescale-prior loss
$L_{\text{smooth}}, L_{\text{bounds}}, L_Q$	Smoothness loss, bounds penalty, and optional source-term regularisation loss
$\lambda_{\text{gw}}, \lambda_{\text{cons}}, \lambda_{\text{prior}}$	Weights assigned to the groundwater, consolidation, and timescale-prior losses

$\lambda_{\text{smooth}}, \lambda_{\text{bounds}}, \lambda_Q$	Weights assigned to the smoothness, bounds, and optional source-term regularisation losses
$\varepsilon_{\text{cons}}$	Root-mean-square consolidation-balance diagnostic
$\varepsilon_{\text{prior}}$	Root-mean-square timescale-prior consistency diagnostic
MAE	Mean absolute error
MSE	Mean squared error
RMSE	Root mean squared error
R^2	Coefficient of determination
PIT	Probability integral transform
BS	Brier score

目 次

摘要	III
Abstract	VII
符号表	IX
插图清单	XIX
附表清单	XX
Chapter 1 Introduction	1
Chapter 2 State of the Art and Scientific Gap	6
2.1 Introduction	6
2.2 Land subsidence as a coupled hydrogeological and urban process	7
2.3 Observation, measurement, and diagnostic analysis	8
2.4 Physically based modelling of subsidence	9
2.5 Data-driven and machine-learning approaches	10
2.6 Toward physics-guided and hybrid frameworks	11
2.7 Scientific gap and positioning of this study	12
2.8 Summary	13
Chapter 3 Study Areas, Data, and Preprocessing	15
3.1 Introduction	15
3.2 Study areas	15
3.3 Data sources and target variables	18
3.4 Cross-city harmonisation framework	20
3.5 Feature construction and input representation	21
3.6 Target definition and forecasting windows	22
3.7 Preprocessing strategy and physical consistency	23
3.8 Scientific rationale of the data framework	24
3.9 Summary	24
Chapter 4 Machine Learning-Based Subsidence Simulation	26
4.1 Introduction	26

4.2	Formulation of the prediction problem	26
4.3	Temporal harmonisation through the delta-rate strategy	27
4.4	Extreme Gradient Boosting Regressor	28
4.4.1	Theory	28
4.4.2	Role in the Nansha study	29
4.5	Long Short-Term Memory network	30
4.5.1	Sequential memory formulation	30
4.5.2	Role in the Nansha study	30
4.6	Scenario-based prevention through modified delta rates	31
4.7	Data normalisation and supervised sequence construction	32
4.8	Calibration and walk-forward validation	32
4.9	Evaluation metrics	33
4.10	Methodological significance of the machine-learning stage	34
4.11	Limitations of the purely data-driven stage	34
4.12	Summary	35
Chapter 5	Physics-Guided AI for Subsidence Forecasting	36
5.1	Introduction	36
5.2	Forecasting task and input representation	39
5.3	General architecture of GeoPriorSubsNet	40
5.4	Predictive backbone: HALNet	42
5.4.1	Probabilistic formulation	42
5.4.2	Embedding and recurrent encoding	42
5.4.3	Attentive fusion and probabilistic decoder	43
5.5	Physics head and effective closure fields	43
5.5.1	Need for a dedicated physics head	43
5.5.2	Drainage thickness and prior-centred fields	44
5.5.3	Closure-consistent relaxation timescale	44
5.5.4	Positivity and admissible bounds	45

5.6	Reduced hydro–geomechanical dynamics	45
5.6.1	Groundwater residual	45
5.6.2	Consolidation residual	46
5.6.3	Residual losses	46
5.7	Probabilistic data loss and composite objective	47
5.7.1	Quantile loss	47
5.7.2	Timescale prior, smoothness, and composite objective	47
5.8	Identifiability and interpretation of the inferred fields	48
5.9	Methodological significance of the physics-guided stage	49
5.10	Summary	49
Chapter 6	Results and Experimental Validation	51
6.1	Introduction	51
6.2	Contrasting driver regimes across Nansha and Zhongshan	51
6.3	Deterministic predictive performance	53
6.3.1	Core out-of-sample accuracy	53
6.3.2	Benchmark comparison against non-physics and physics-based baselines	55
6.4	Probabilistic calibration, sharpness, and horizon stability	55
6.5	Forecasted hotspot evolution and persistence	58
6.6	Inferred effective fields and physics diagnostics	59
6.6.1	Closure consistency and dynamic balance	59
6.6.2	Spatial structure of the inferred fields	60
6.6.3	Identifiability and remaining trade-offs	61
6.7	Point-support external validation	63
6.8	Cross-city transfer and deployability	64
6.8.1	Retention of deterministic skill	64
6.8.2	Transfer of uncertainty and threshold risk	65
6.8.3	Transfer of hotspot ranking	66
6.9	Summary of the experimental evidence	67

Chapter 7 Discussion	68
7.1 Introduction.....	68
7.2 From data-driven simulation to physics-guided forecasting	68
7.3 Scientific interpretation of the two urban subsidence regimes.....	70
7.4 Methodological significance of the physics-guided framework	71
7.5 Interpretability, identifiability, and the meaning of the inferred fields.....	73
7.6 Practical significance for urban risk management	74
7.7 Limitations of the study	75
7.8 Future research directions.....	76
7.9 Conclusion and perspectives.....	77
References	79
Achievements	89
Acknowledgements	91

插图清单

3-1	Field evidence of differential settlement in Nansha and Zhongshan. a , Nansha: oblique satellite view with locations of selected InSAR pixels (NS5111–1, NS5101–1, NS5189–1) and photographs of nearby buildings. Yellow arrows indicate vertical offsets between building plinths and the surrounding ground, with measured gaps of approximately 402–441 mm at several riverfront structures. b , Zhongshan: similar layout for pixels ZJ1921–1, ZJ1922, ZJ1925 and ZJ1940–1, illustrating tilted walls, uplifted steps, and exposed foundations with offsets of roughly 105–342 mm. Insets show the wider building context. These field observations corroborate the persistent subsidence hotspots identified by the InSAR time series in both cities.	17
5-1	Overall workflow of GeoPriorSubsNet. The framework integrates heterogeneous data sources, harmonization and preprocessing, construction of static and dynamic predictor streams, future-known exogenous forcing, and a dedicated physics input through H_{eff} . These components feed the GeoPriorSubsNet architecture, which combines a probabilistic data head and a physics head within a joint training objective. The workflow further includes post-hoc calibration, physical diagnostics, cross-city transfer tests, and multi-horizon forecasting outputs.	38

5-2	Overview of the GeoPriorSubsNet forecasting architecture. The model combines a hybrid attentive recurrent predictive backbone, denoted HALNet, with a physics head constrained by reduced hydro–geomechanical dynamics. Static, dynamic, and future-known inputs are embedded and fused into horizon-specific latent states. These latent states feed both probabilistic forecast heads for subsidence and hydraulic head, and a physics head that infers effective hydraulic conductivity, specific storage, and closure-consistent relaxation timescale. The framework therefore couples predictive learning, uncertainty quantification, and physical consistency within a single architecture.	41
6-1	Driver–response relationships between cumulative subsidence and the main controlling variables in Nansha and Zhongshan. The panels show the empirical relation between cumulative subsidence and groundwater depth, rainfall, effective soil thickness, and shallow urban-load proxy. The black curves indicate the smoothed trend. Nansha exhibits broader hydro-environmental variability and stronger monotonic sensitivity to groundwater depth and thickness, whereas Zhongshan shows a shallower groundwater regime and a more compact driver space.	53
6-2	Core deterministic performance and no-physics ablation. The figure summarises the out-of-sample performance of GeoPriorSubsNet in Nansha and Zhongshan, together with the effect of removing the physics constraints. The physics-guided configuration retains strong deterministic skill in both cities and yields a particularly clear advantage in Zhongshan, where it improves both accuracy and uncertainty sharpness relative to the no-physics alternative.	54

6-3	Reliability, horizon-wise calibration, and sharpness. Panels a–b show the reliability curves for Nansha and Zhongshan. Panel c shows the horizon-specific reliability behaviour from H1 to H3. Panel d summarises the horizon-wise sharpness. The results indicate that the predictive distribution remains reasonably well calibrated across horizons, with broader intervals in Nansha and sharper but still well-targeted intervals in Zhongshan.	57
6-4	Additional probabilistic diagnostics. The figure reports PIT histograms for Nansha and Zhongshan, horizon-wise coverage against the nominal 80% target, sharpness evolution with lead time, calibration factors per horizon, and quantile residuals across horizons. Together these diagnostics confirm that uncertainty behaviour remains stable with lead time, although Nansha requires broader intervals than Zhongshan.	58
6-5	Forecasted hotspot evolution and persistence in Nansha and Zhongshan. Panels a–b and f–g show the forecasted annual subsidence change and exceedance probability maps. Panels c–d and h–i summarise hotspot counts, temporal evolution, and priority clusters. Panels e and j show hotspot persistence. The results indicate broader and more persistent risk corridors in Nansha and more localised hotspot organisation in Zhongshan.	59
6-6	Distribution of the consolidation residual R_{cons} in Nansha and Zhongshan. In both cities, the residual distribution is sharply concentrated near zero, consistent with the very small values of $\varepsilon_{\text{cons}}$. This indicates that the forecast trajectories remain closely aligned with the reduced relaxation law imposed by the physics-guided framework.	60

6-7	Synthetic identifiability experiments. Panel a shows timescale recovery, panel b permeability recovery through the closure, panel c the degeneracy-ridge diagnostic, and panel d the relation between error and the identifiability metric. The results show that the closure-consistent timescale is more robustly recoverable than the exact decomposition into conductivity, storage, and drainage thickness.	62
6-8	Bounds and ridge summary across synthetic realisations. The figure reports bound hits, the distribution of ridge non-identifiability, the relation between clipping and ridge behaviour, and the lithology-dependent distribution of failure modes. Strong ridge behaviour occurs only in a minority of realisations, whereas most solutions remain away from severe non-identifiability.	62
6-9	Point-support external validation in Zhongshan. Panel a shows the H_{eff} field and the locations of the five validation sites. Panel b compares borehole-derived compressible thickness with model H_{eff} . Panel c compares late-step specific capacity with model conductivity K . The results show a meaningful thickness-pathway agreement but a weaker relation between local pumping productivity and grid-scale effective conductivity.	64
6-10	Cross-city transfer performance. The figure reports retention of deterministic skill, horizon-wise RMSE retention, uncertainty behaviour under transfer and warm-start adaptation, reliability for the 50 mm yr^{-1} exceedance event, Brier score, Jaccard overlap of the top-ranked hotspots, and Spearman rank correlation of hotspot ordering. Direct zero-shot transfer is unreliable, whereas warm-start adaptation restores most of the target-city skill and substantially improves uncertainty quality.	66

附表清单

4-1	Hyperparameter configuration used for the machine-learning subsidence simulation in Nansha.	33
-----	---	----

Chapter 1 Introduction

Land subsidence is one of the most consequential forms of ground deformation affecting urban environments. It is expressed as the progressive lowering of the land surface under the combined action of natural processes and human activities, and it can induce long-lasting damage to buildings, transport corridors, buried utilities, drainage systems, and flood-protection infrastructure [1-5]. Although the phenomenon has been recognised for a long time [6], its present significance is much greater because many of the fastest-growing cities are expanding precisely in sedimentary lowlands, deltaic plains, and coastal regions where the subsurface is highly susceptible to compaction [7-10]. In such settings, land subsidence is no longer a localised geotechnical inconvenience; it has become a strategic urban, hydrological, and environmental problem.

The urgency of this problem arises from the convergence of several pressures. Rapid urbanisation increases structural loading, modifies surface runoff and recharge pathways, and intensifies groundwater demand. At the same time, subsidence reduces freeboard, degrades drainage efficiency, and amplifies exposure to pluvial, fluvial, and coastal flooding. In low-lying coastal zones, even moderate subsidence can substantially increase relative sea-level rise and therefore alter the scale of future flood hazard [8-9, 11]. Recent global assessments have further shown that subsidence is not restricted to a few emblematic cities, but affects a wide range of urban and agricultural regions, with major implications for groundwater storage loss, infrastructure stability, and long-term land-use planning [7, 10]. The problem is therefore both scientifically challenging and societally urgent.

From a process perspective, land subsidence is intrinsically multifactorial. In many basins, excessive groundwater withdrawal is the dominant driver because it lowers pore pressure, increases effective stress, and promotes the irreversible compaction of compressible sediments such as clays and silts [12-14, 3, 15]. However, groundwater depletion alone does not explain the full diversity of observed deformation patterns. Subsidence

is also shaped by lithological architecture, aquifer–aquitard connectivity, drainage conditions, natural consolidation history, tectonic setting, and anthropogenic loading associated with buildings, transport infrastructure, land reclamation, and urban surface sealing [16, 1, 17-19]. The resulting deformation is therefore the macroscopic outcome of coupled hydrogeological, geomechanical, and urban-development processes operating across multiple spatial and temporal scales.

This complexity explains why reliable forecasting remains difficult. Traditional physically based approaches, including groundwater-flow, poroelastic, and consolidation models, provide an essential process-oriented framework and remain indispensable for mechanism-based interpretation [14, 20-23]. Yet their deployment at urban scale is often constrained by the need for dense hydrogeological parameterisation, boundary-condition specification, and repeated high-cost simulation. In many rapidly developing cities, such detailed subsurface characterisation is either incomplete or unavailable. Conversely, purely empirical and machine-learning approaches can discover strong predictive relationships from heterogeneous geospatial data, but they often remain weak in physical interpretability and may produce forecasts that are statistically plausible without being mechanistically credible [24-28]. The central methodological problem is therefore clear: urban subsidence forecasting requires models that are both predictive and physically accountable.

This work addresses that problem through two complementary research stages. The first stage established a robust data-driven basis for land subsidence simulation in a rapidly urbanising environment. In that study, advanced machine-learning models, notably the eXtreme Gradient Boosting Regressor and the Long Short-Term Memory network, were used to analyse and forecast subsidence in the Nansha district. The results showed that groundwater level and building concentration were among the most influential explanatory variables, and that well-designed machine-learning models could reproduce the observed deformation dynamics with strong predictive skill [29]. That contribution was important because it showed

that heterogeneous urban and environmental data could be integrated into an operational forecasting workflow and that machine learning can support risk-oriented simulation in settings where conventional modeling remains difficult to calibrate.

However, the first stage also made a major limitation explicit. Even when predictive performance is strong, a data-driven model does not, by itself, guaranty physical consistency. It may capture correlations that are highly useful for interpolation within the observed range but remain difficult to interpret in terms of groundwater depletion, storage, drainage conditions, or compaction timescales. This limitation becomes more critical when the purpose of modeling is not only to fit past deformation, but also to support extrapolation, uncertainty analysis, scenario testing, and decision-making in environments subject to non-stationary forcing. In urban hazard forecasting, predictive accuracy alone is not sufficient; forecasts must also remain consistent with the governing physical structure of the problem.

The second stage of this research therefore moves beyond black-box forecasting and develops a physics-guided artificial intelligence framework for land subsidence prediction. The central idea is not to replace physical reasoning with machine learning, nor to force a full numerical hydrogeological model where the required inputs are unavailable, but to construct an intermediate framework in which deep learning and reduced physical structure inform one another. In this study, that framework is represented by a geomechanically constrained forecasting architecture designed to couple heterogeneous spatio-temporal inputs with physically interpretable effective fields, including hydraulic conductivity, specific storage, and consolidation timescale, while producing calibrated multi-horizon forecasts at urban scale. Such a direction is consistent with the broader evolution of physics-informed learning, which seeks to incorporate governing constraints into data-driven models in order to improve plausibility, generalisation, and interpretability [30-33].

The scientific novelty of this second stage lies in the fact that the model does not treat physics as a loose regularisation slogan. Instead, the physical

structure is formulated so that the learned dynamics can be audited against a reduced consolidation balance and against a closure-consistent relationship between hydrogeological fields and emergent relaxation timescale. This is a decisive step for subsidence forecasting. It allows the model to remain compatible with the resolution and diversity of urban geospatial data, while also providing mechanistic diagnostics that can distinguish a purely statistical fit from a physically coherent forecast. In practical terms, this means that the model is designed not only to estimate future deformation, but also to explain how inferred subsurface properties and system timescales organise the predicted response across contrasting urban regimes.

A second element of novelty concerns uncertainty. Many previous subsidence studies have concentrated primarily on deterministic prediction or susceptibility mapping [24-25, 34]. Yet urban decision-making rarely depends on a single expected value alone. It depends on the probability that a threshold will be exceeded, on the stability of hotspot ranking across forecast horizons, and on the ability to distinguish persistent high-risk zones from more transient ones. The physics-guided framework developed in this work therefore targets probabilistic, multi-horizon forecasting rather than only point-wise prediction. This perspective is closely aligned with recent progress in interpretable time-series forecasting [35-36] and is particularly relevant for subsidence-prone cities where planning decisions must be made under uncertainty.

A third contribution is comparative. Rather than limiting the analysis to a single urban setting, the research examines contrasting subsidence regimes represented by Nansha and Zhongshan. This comparison is scientifically important because it allows the forecasting framework to be tested across different hydrostratigraphic and urban-development contexts. Nansha represents a rapidly evolving coastal and reclaimed environment where subsidence is closely tied to strong anthropogenic forcing and complex shallow sedimentary conditions. Zhongshan, by contrast, provides a more structured setting with different lithological and hydrogeological controls. Investigating both regimes improves the interpretability of the results and,

address not only predictive accuracy, but also physical consistency, transferability, and the limits of generalisation under distribution shift.

The broader impact of this research is both methodological and applied. Methodologically, it contributes to the ongoing effort to build forecasting models that are more transparent, more interpretable, and more physically grounded than conventional black-box systems, while remaining lighter and more scalable than fully parameterised numerical simulators. Applied-wise, it supports a more proactive view of urban subsidence management. Forecasts that are spatially resolved, physically interpretable, and uncertainty-aware can help identify where intervention is most urgent, where monitoring should be strengthened, and where groundwater management or construction control may yield the greatest benefit. This is particularly relevant in the context of sustainable urban development, groundwater protection, and climate adaptation, and therefore speaks directly to the broader agenda of resilient cities and responsible resource governance [37-38].

Against this background, the present report is organised around one central scientific question: how can urban land subsidence forecasting be advanced from accurate data-driven simulation toward physically guided, interpretable, and decision-relevant prediction? The report addresses this question by first establishing the geohazard context and the methodological gap, then presenting the data foundation and study areas, then developing the two successive modelling stages, and finally analysing the experimental results, their physical meaning, and their implications for urban risk governance. In this sense, the report is not a juxtaposition of two independent studies. It documents a research progression: from machine-learning-based subsidence simulation, which demonstrated the predictive value of heterogeneous urban data, to physics-guided artificial intelligence, which seeks to restore mechanistic structure, uncertainty awareness, and scientific accountability to urban subsidence forecasting.

Chapter 2 State of the Art and Scientific Gap

2.1 Introduction

Land subsidence is a progressive lowering of the ground surface caused by natural and anthropogenic processes acting on deformable subsurface materials. Although the phenomenon has long been recognized in the geosciences [6], its present importance is much greater because it now affects rapidly urbanizing, groundwater-stressed, and coastal regions where infrastructure, population, and economic assets are highly concentrated [7-10]. In such environments, subsidence is not simply a geological phenomenon. It is a major urban geohazard that may damage buildings, roads, railways, buried pipelines, drainage systems, and flood-control infrastructure, while also amplifying pluvial and coastal flood exposure [2, 39, 11].

The scientific challenge of subsidence forecasting arises from the fact that the observed deformation is the outcome of coupled hydrogeological, geomechanical, geomorphological, and urban-development processes. Groundwater withdrawal is frequently the dominant driver, but its effect is modulated by lithology, aquifer–aquitard architecture, delayed drainage, anthropogenic loading, land reclamation, and local geological history [12-13, 17, 19]. As a result, the same deformation signal can emerge from different combinations of processes and material properties, which makes interpretation and prediction difficult.

This chapter reviews the scientific foundations of urban subsidence forecasting and identifies the gap that motivates the present study. The discussion begins with the physical and urban controls of land subsidence, then examines the main observational and modelling approaches, including InSAR-based monitoring, process-based numerical simulation, and machine-learning forecasting. The chapter finally explains why a new framework is needed, namely one that combines predictive performance, physical interpretability, and uncertainty-aware forecasting within a single scalable model.

2.2 Land subsidence as a coupled hydrogeological and urban process

Land subsidence has been documented in many geological and urban settings, including alluvial basins, coastal plains, deltaic environments, reclaimed lands, mining districts, and intensively pumped aquifer systems [3, 16, 7-8]. In urban regions, its consequences are amplified because the deformation interacts with concentrated infrastructure and engineered surfaces. Even moderate differential settlement may induce cracking of buildings, distortion of transport networks, malfunction of buried utilities, and increased flood susceptibility in low-lying areas [1-2, 9].

In many sedimentary basins, groundwater extraction is the principal mechanism of long-term subsidence. Declining hydraulic head reduces pore pressure and increases effective stress, promoting the compaction of compressible sediments such as clays and silts [12-14]. In its most classical form, this relationship can be expressed by Terzaghi's effective stress principle,

$$\sigma' = \sigma - p, \quad (2-1)$$

where σ' is the effective stress, σ is the total stress, and p is pore pressure. When groundwater withdrawal lowers p , the effective stress increases, and compaction may occur if the sediments are sufficiently compressible.

However, groundwater decline alone does not determine the final subsidence pattern. The spatial and temporal response depends strongly on lithological composition, thickness of compressible units, drainage path length, aquifer connectivity, and delayed consolidation effects [15, 40, 18]. Additional anthropogenic influences, such as urban loading, land reclamation, and engineering dewatering, may further modify the deformation field [16-17, 19]. Tectonic or mining-related disturbance can also interact with subsidence in certain settings [4, 41]. For this reason, subsidence must be treated as a coupled system rather than a single-cause phenomenon.

Two implications follow from this complexity. First, a realistic forecasting framework must combine hydrogeological, geomechanical, and urban information rather than relying on a single driver. Second, the relative

importance of these controls is basin-dependent. A model that performs well in one city may not generalize to another if it is not sufficiently sensitive to the underlying physical and geological context.

2.3 Observation, measurement, and diagnostic analysis

The observation of land subsidence has been transformed by satellite geodesy, especially Interferometric Synthetic Aperture Radar (InSAR). Differential InSAR and time-series InSAR techniques allow ground deformation to be mapped at millimetric precision over wide urban areas, providing a major advantage over point-based monitoring methods such as leveling or isolated GNSS stations [42-45]. These methods have revealed persistent deformation hotspots in many cities, including Mexico City, Beijing, Wuhan, Shenzhen, Shanghai, and Semarang [1, 46-50].

The observational literature has established two major facts. First, subsidence is often highly heterogeneous in space, even within the same urban basin. Sharp gradients may arise over short distances because of contrasts in sediment thickness, groundwater exploitation, or construction intensity. Second, subsidence is temporally non-stationary. Long-term trends may be modulated by hydrological recovery, pumping cycles, engineering projects, or management interventions [51-53]. These findings have led to the increasing integration of InSAR with groundwater records, borehole logs, pumping tests, land-use information, and infrastructure data in order to move from deformation mapping toward process interpretation [18, 54-55].

Despite these advances, observation alone does not solve the forecasting problem. InSAR provides a high-resolution record of surface displacement, but it does not directly resolve the subsurface state variables that govern deformation. A deformation map is therefore a diagnostic product rather than a predictive model. The scientific task is to move from observation to dynamics, that is, from measured surface motion to a model able to explain past subsidence and credibly forecast future evolution.

2.4 Physically based modelling of subsidence

Physically based models remain the main framework for representing the causal chain linking groundwater flow, effective stress change, and compaction. In general form, groundwater flow in a deformable medium may be written as

$$S_s \frac{\partial h}{\partial t} - \nabla \cdot (K \nabla h) = Q, \quad (2-2)$$

where S_s is specific storage, h is hydraulic head, K is hydraulic conductivity, and Q is a source or sink term. This equation expresses how head evolves under storage, flow, and pumping or recharge forcing. Compaction may then be related to head decline and to the mechanical response of compressible units through poroelastic or consolidation formulations [14, 13, 56].

A large body of work has shown the value of these models in regional or basin-scale applications. They have been used to reconstruct the effects of groundwater exploitation, simulate delayed compaction, and test mitigation strategies in Beijing, Tehran, and other subsiding basins [57, 40-41]. More recent work has also introduced data assimilation to better anchor model states and parameters to observed deformation fields [20-22, 58, 23]. These studies represent a significant advance because they reduce the disconnect between process-based simulation and surface observations.

Nevertheless, physically based modelling remains difficult in urban forecasting practice. Reliable application usually requires dense hydrostratigraphic characterization, site-specific mechanical and hydraulic parameters, well-constrained boundary conditions, and long time series of groundwater and deformation observations. In many rapidly urbanizing settings, such information is incomplete, spatially uneven, or unavailable at the resolution needed for robust inversion. As a consequence, fully process-based models may become underconstrained or computationally expensive, particularly when extended to wide urban domains or repeated scenario analysis. There is therefore a persistent trade-off between physical realism and practical scalability.

2.5 Data-driven and machine-learning approaches

Machine learning has become increasingly important in subsidence research because it can exploit high-dimensional, nonlinear, and heterogeneous geospatial datasets without requiring a fully specified physical model. Tree-based algorithms, support vector machines, geographically weighted deep learning, recurrent neural networks, and ensemble methods have been used to map susceptibility, identify important drivers, and forecast deformation [24, 59-60, 34, 25, 27, 26, 61]. In a general forecasting setting, these approaches aim to approximate a nonlinear operator

$$\hat{s}_{t+h} = f_{\theta}(x_{1:t}, s_{1:t}), \quad (2-3)$$

where \hat{s}_{t+h} is future subsidence, $x_{1:t}$ denotes the history of explanatory variables, $s_{1:t}$ denotes past deformation, and f_{θ} is a data-driven model parameterized by θ .

The literature has shown that machine-learning models can recover meaningful predictive relationships. Studies in Beijing and other regions have identified groundwater decline, sediment thickness, building density, and human activity indicators as important controls [24, 62, 18]. Recent work has extended these methods to spatio-temporal modelling and scenario-based simulation [25, 27, 26]. In the context of the present study, Liu et al. [29] showed that machine-learning models can simulate subsidence in a rapidly urbanizing environment with strong predictive skill while identifying groundwater level and building concentration as key controlling variables.

However, the literature also shows clear limitations. Many machine-learning studies focus on susceptibility mapping or retrospective fitting rather than true multi-horizon forecasting. Even when predictive accuracy is high, the internal representation of the model often remains difficult to interpret in hydrogeological or geomechanical terms. This is a serious issue because subsidence is not merely a statistical classification problem. It is a dynamic physical response of a deformable urban subsurface. A purely empirical model may interpolate within the training domain ef-

fectively while remaining weak under extrapolation, distribution shift, or policy-driven scenario changes. As emphasized by Guzy and Malinowska [28], the challenge is not only to predict where subsidence occurs, but to explain why it occurs and how it may evolve under changing forcing.

2.6 Toward physics-guided and hybrid frameworks

The limitations of both purely physical and purely data-driven models have motivated a new line of research based on hybrid or physics-guided learning. In this setting, machine learning is combined with reduced governing equations, physical priors, or structured constraints in order to improve plausibility, regularity, and generalization [30-33]. The idea is not to replace physics with data, nor to impose a full numerical simulator where the required inputs are unavailable. Rather, it is to construct a model that learns flexibly from data while remaining anchored to core physical principles.

For land subsidence, this direction is particularly relevant. The available observations are often rich enough to support data-driven learning, but not sufficient to sustain a fully resolved hydrogeological inversion at urban scale. A reduced but physically meaningful closure can therefore provide an attractive compromise. Instead of requiring all hydrostratigraphic details explicitly, such a framework may constrain the forecast through effective storage, conductivity, drainage thickness, or consolidation timescale. Some recent data-assimilation studies already point toward this philosophy by combining observations with physically structured representations rather than treating the deformation field as a purely empirical surface [20-23].

In reduced form, a physically guided subsidence dynamics may be expressed through a relaxation law such as

$$\frac{\partial s}{\partial t} = \frac{s_{\text{eq}} - s}{\tau}, \quad (2-4)$$

where s is cumulative subsidence, s_{eq} is an equilibrium settlement linked to drawdown and compressible thickness, and τ is an effective relaxation timescale. This type of formulation does not replace full poroelastic mod-

elling, but it captures an essential aspect of delayed compaction behavior while remaining compatible with the scale and incompleteness of urban datasets.

Even so, a major gap remains in the literature. Existing studies often belong to one of three broad groups. The first provides high-quality observation and diagnosis but stops short of predictive mechanistic modelling. The second applies process-based simulation with strong physical fidelity but substantial parameter and calibration demands. The third uses machine learning for prediction or susceptibility mapping, often with limited physical interpretability and limited treatment of forecast uncertainty. What remains uncommon is a single framework that combines urban-scale predictive skill, physically auditable internal structure, and uncertainty-aware forecasting.

2.7 Scientific gap and positioning of this study

The scientific gap addressed in this research can therefore be stated clearly. Urban land subsidence studies have achieved major progress in observation, process-based modelling, and data-driven prediction, but these advances remain only partially integrated. Observation reveals where deformation occurs. Physics explains how compaction can arise under specified hydraulic and mechanical conditions. Machine learning improves predictive efficiency from heterogeneous data. The unresolved challenge is to combine these strengths in a forecasting framework that is sufficiently accurate for urban-scale application, sufficiently structured for physical interpretation, and sufficiently flexible for complex multi-source datasets.

The present study positions itself directly within that gap. A first contribution is to establish a rigorous machine-learning basis for subsidence simulation in a real urban environment, showing that heterogeneous geospatial predictors can yield accurate deformation estimates and meaningful driver analysis [29]. A second contribution is to move beyond black-box prediction by introducing a physics-guided artificial-intelligence framework in which forecasts are informed by reduced consolidation dy-

namics and by interpretable effective hydrogeological fields. In schematic form, the target is to learn a mapping of the form

$$\left(\hat{s}_q, \hat{h}_q, K, S_s, \tau\right) = \mathcal{M}_\Theta(X), \quad (2-5)$$

where \hat{s}_q and \hat{h}_q denote forecast quantiles for subsidence and head, K is effective hydraulic conductivity, S_s is effective specific storage, τ is effective relaxation timescale, and \mathcal{M}_Θ is a physics-guided model learned from input data X .

The novelty of this research lies at the intersection of three requirements that are rarely satisfied simultaneously in the existing literature. The first is predictive capability at urban scale from heterogeneous spatio-temporal inputs. The second is physical accountability, meaning that the learned dynamics can be interpreted and audited in terms of subsurface response rather than only statistical fit. The third is uncertainty-aware forecasting, which is essential for decision-making in risk-prone urban systems. The scientific progression developed in this work is therefore deliberate: the machine-learning study establishes what can be learned directly from the available data, and the physics-guided study addresses what remains missing when prediction is detached from mechanism.

2.8 Summary

This chapter reviewed the state of the art in urban land subsidence research from the perspectives of physical mechanisms, observational diagnostics, process-based modelling, machine-learning prediction, and physics-guided hybrid frameworks. The review shows that land subsidence is a coupled hydrogeological and urban process, whose forecasting remains difficult because the deformation signal arises from multiple drivers acting across different spatial and temporal scales. InSAR and related geodetic methods have greatly improved observation, but they do not by themselves provide a predictive dynamical model. Physically based approaches offer strong process fidelity, yet often require more information than is available in rapidly urbanizing basins. Data-driven approaches offer flexibility and predictive power, but often lack mechanistic interpretability and robust

uncertainty treatment.

The literature therefore reveals a clear scientific gap. What is still needed is a forecasting framework that combines urban-scale predictive skill, physically interpretable effective structure, and uncertainty-aware prediction. This gap motivates the methodological development of the following chapters. The next chapter presents the study areas, data foundation, and preprocessing strategy used to build such a framework.

Chapter 3 Study Areas, Data, and Preprocessing

3.1 Introduction

This study is built on a multi-source urban subsidence dataset designed to support both data-driven simulation and physics-guided forecasting. The objective is not only to assemble deformation and environmental observations, but to organize them into a coherent modelling framework that preserves physical interpretability across contrasting urban basins. The data structure must therefore satisfy two requirements. First, it must be rich enough to support nonlinear spatio-temporal forecasting. Second, it must remain sufficiently structured to allow hydrogeological and geomechanical interpretation.

Two subsidence-prone urban environments in the Pearl River Delta of southern China are considered: Nansha District, Guangzhou, and Zhongshan. Their joint use is deliberate. It allows the analysis to move beyond a single-city demonstration and to examine forecasting under contrasting geological, hydrogeological, and urban-development conditions. In the first stage of the study, Nansha served as the basis for machine-learning-based subsidence simulation [29]. In the second stage, both Nansha and Zhongshan were treated under a common city-agnostic framework in order to assess physical consistency, uncertainty quality, and cross-city transferability.

3.2 Study areas

Nansha District is located at the southernmost part of Guangzhou, along the western sector of the Pearl River Delta waterway. It occupies a strategic position near the confluence of the Xijiang, Beijiang, and Dongjiang rivers and has undergone intense urbanization during recent decades [29]. The district is characterized by a low-lying deltaic setting, widespread Quaternary sediments, and strong human modification associated with rapid construction, land development, and growing pressure on groundwater and infrastructure. Such conditions make Nansha espe-

cially relevant for subsidence research, because deformation in coastal and recently reclaimed plains often arises from the combined action of compressible deposits, groundwater stress, and anthropogenic loading [16, 19, 8].

Zhongshan provides a contrasting urban basin within the same regional context. In this work, it is analysed on the same annual grid and under the same harmonisation rules as Nansha. Compared with Nansha, Zhongshan shows a more competent lithological composition, a shallower groundwater-depth range, and a more spatially structured deformation regime. By contrast, Nansha contains a larger fraction of fine-grained, compressible materials and exhibits a broader hydro-environmental span [18, 48]. This contrast is scientifically important because it allows the forecasting framework to be tested under two different classes of urban subsidence behaviour: a highly heterogeneous, rapidly reclaimed coastal system, and a more organised urban basin with clearer hydrostratigraphic structure.



图 3-1 Field evidence of differential settlement in Nansha and Zhongshan. **a**, Nansha: oblique satellite view with locations of selected InSAR pixels (NS5111-1, NS5101-1, NS5189-1) and photographs of nearby buildings. Yellow arrows indicate vertical offsets between building plinths and the surrounding ground, with measured gaps of approximately 402–441 mm at several riverfront structures. **b**, Zhongshan: similar layout for pixels ZJ1921-1, ZJ1922, ZJ1925 and ZJ1940-1, illustrating tilted walls, uplifted steps, and exposed foundations with offsets of roughly 105–342 mm. Insets show the wider building context. These field observations corroborate the persistent subsidence hotspots identified by the InSAR time series in both cities.

The combined use of Nansha and Zhongshan serves three purposes. First, it grounds the analysis in real urban case studies of strong practical relevance. Second, it provides a basis for comparing driver–response behaviour across cities rather than attributing all observed patterns to a single local setting. Third, it supports the transition from machine-learning simulation to physics-guided forecasting, where generalisation, mechanistic consistency, and uncertainty calibration become central concerns.

3.3 Data sources and target variables

The data foundation of this study is built from heterogeneous geospatial, hydro-environmental, and geological observations assembled to represent the principal controls of urban land subsidence. In the machine-learning stage, the Nansha dataset integrated InSAR-derived deformation, geological information, filled-soil thickness, sand-layer distribution, building concentration, and groundwater levels [29]. In the later physics-guided framework, this data contract was reformulated so that similar classes of information could be harmonised across Nansha and Zhongshan on a common annual grid. The resulting archive includes InSAR-derived vertical displacement, groundwater observations, rainfall fields, borehole lithology and thickness information, and urban-development layers such as imperviousness, roads, and low-rise building footprints [42, 1, 26, 47].

The principal target variable throughout the study is cumulative land subsidence, denoted here by $s(i, t)$ for grid cell i and year t . This quantity is derived from multi-temporal InSAR processing and is used because it summarizes the integrated ground response over time, while remaining directly relevant to urban hazard and infrastructure risk. In the machine-learning framework, the Nansha deformation archive covered the period 2015–2022 and served as the basis for simulation and scenario analysis [29]. In the later framework, the cumulative subsidence target was harmonised on an annual support for both Nansha and Zhongshan, together with groundwater-state information where available.

Groundwater information forms the second major data component.

The literature has consistently shown that groundwater decline is one of the most powerful controls on urban subsidence [12-13, 15, 24-25]. This is also consistent with the present study, where groundwater depth emerged as a dominant driver in Nansha and remained central in the physics-guided formulation. Let $d(i, t)$ denote depth to water, measured positive downward from the surface. The corresponding hydraulic head is reconstructed as

$$h(i, t) = z_{\text{surf}}(i) - d(i, t), \quad (3-1)$$

where $z_{\text{surf}}(i)$ is the surface elevation datum. The associated drawdown is then written as

$$\Delta h(i, t) = h_{\text{ref}}(i) - h(i, t), \quad (3-2)$$

where $\Delta h > 0$ indicates head loss relative to a reference state. This canonicalisation is essential because it preserves a consistent sign convention between data ingestion and physical interpretation.

Rainfall is included as an additional dynamic forcing variable. Although its influence is generally weaker than that of groundwater withdrawal, rainfall affects recharge, hydraulic recovery, and the interannual modulation of subsurface stress conditions [51-52]. It therefore contributes both to predictive completeness and to a more realistic hydrogeological interpretation. In the physics-guided framework, rainfall is used both as an observed historical driver and as a future-known exogenous input over the forecast horizon.

Subsurface geological information is represented through lithology and thickness-related variables. In Nansha, the first modelling stage used lithological descriptors, filled-soil thickness, and sand-layer distribution [29]. In the later framework, these components were reorganised into a thickness-aware and lithology-aware representation suited to cross-city learning. Borehole information was used to derive an effective compressible-thickness field, $H_{\text{eff}}(i)$, while lithological structure was preserved as a static descriptor. This variable is scientifically important because both the literature and the Nansha results show that compressible-layer thickness is a major amplifier of deformation [62, 40, 25]. In

Zhongshan, the available thickness information includes right-censoring, which was treated explicitly during preprocessing rather than ignored.

Urban-development information was introduced to represent the anthropogenic stress imposed by built environments. In the earlier machine-learning stage, building concentration was used directly and was identified as one of the most influential predictors of deformation in Nansha [29]. In the later framework, this information was generalised into a shallow urban-load proxy, denoted $U^*(i, t)$, derived from urban surface characteristics such as imperviousness and built-form patterns. Its role is not to model the full mechanics of foundations and structures, but to capture the distributed shallow loading and surface modification associated with intense urbanisation.

3.4 Cross-city harmonisation framework

A major objective of this study was to move from a single-city modelling exercise to a common forecasting framework that permits direct comparison across cities. This required strict harmonisation of all input layers. Nansha and Zhongshan were therefore analysed on a common spatial grid and annual time axis, with the same preprocessing rules applied to both study areas. This city-agnostic pipeline ensures that cross-city comparisons and transfer experiments are not confounded by incompatible definitions or inconsistent temporal support.

The harmonisation begins with co-registration of all available layers onto a shared analysis grid. InSAR deformation, groundwater observations, rainfall fields, borehole-derived variables, and urban-development layers are spatially aligned so that each grid cell stores a consistent set of predictors and targets for a given year. Temporal harmonisation then aggregates the records to annual resolution. This annual support is appropriate because the objective is multi-year urban subsidence forecasting rather than short-period fluctuation tracking, while still preserving enough temporal structure for delayed response analysis.

Quality control and masking were applied before model construction.

Pixels or records affected by insufficient support, inconsistent alignment, or invalid values were excluded from the final analysis domain in order to prevent artefacts from propagating into the learning stage. This is particularly important in InSAR-based studies, where decorrelation, intermittent coverage, or local noise may degrade certain time–space locations even when the broader deformation pattern remains reliable [42, 45, 47]. The harmonised archive was therefore designed not as a simple overlay of layers, but as a curated modelling dataset with a consistent physical meaning.

3.5 Feature construction and input representation

After harmonisation, the raw observations were reorganised into feature sets suitable for both machine-learning and physics-guided forecasting. This stage is central because it defines the interface between geoscientific observations and predictive modelling.

In the first stage of the study, feature construction in Nansha focused on assembling the variables most directly linked to subsidence risk, including geology, groundwater level, building concentration, and thickness-related indicators, then learning their relationship to the InSAR-derived target through XGBR and LSTM models [29, 63-64]. A delta-rate strategy was introduced there to strengthen temporal simulation and to make the learning task more responsive to evolving subsidence conditions.

In the second stage, feature construction became more structured. The harmonised archive was packaged into three streams:

$$X_{i,t} = \left(x_i^{(\text{sta})}, X_{i,t}^{(\text{dyn})}, X_{i,t}^{(\text{fut})} \right), \quad (3-3)$$

where $x_i^{(\text{sta})}$ denotes static features, $X_{i,t}^{(\text{dyn})}$ denotes the observed historical sequence, and $X_{i,t}^{(\text{fut})}$ denotes future-known exogenous inputs.

Static features encode time-invariant site properties such as lithology and surface elevation. Dynamic features encode the observed history over the look-back window and include groundwater depth, past cumulative subsidence as an autoregressive channel, rainfall, the shallow urban-load proxy, and a thickness-censor indicator. Future-known features encode

exogenous information available over the forecast horizon; in the present configuration this consists of rainfall only [35-36]. This separation between static, dynamic, and future-known inputs is methodologically important because it clarifies what the model learns from site structure, what it learns from past system evolution, and what it is allowed to condition on at forecast time.

Two engineered variables deserve particular emphasis. The first is the effective compressible-thickness field H_{eff} , which condenses borehole and stratigraphic information into a quantity more directly aligned with consolidation behaviour. The second is the shallow urban-load proxy U^* , which summarises anthropogenic surface pressure in a form usable by both predictive and physical pathways. These variables are central to the methodological development of later chapters because they connect observational data with simplified geomechanical reasoning.

3.6 Target definition and forecasting windows

The principal target is cumulative land subsidence at grid-cell scale. For each grid cell i and forecast origin t , the future target path over a horizon of length H may be written as

$$y_{i,t}^{(s)} = [s(i, t + 1), s(i, t + 2), \dots, s(i, t + H)]. \quad (3-4)$$

In the machine-learning stage, this target is used to simulate past evolution and generate future subsidence scenarios for Nansha. In the physics-guided stage, it is used for probabilistic multi-horizon forecasting in both Nansha and Zhongshan.

Where available, groundwater-state targets are also incorporated:

$$y_{i,t}^{(h)} = [h(i, t + 1), h(i, t + 2), \dots, h(i, t + H)]. \quad (3-5)$$

This dual-target structure is important because the physical residuals linking consolidation and hydraulic response require a coherent representation of both deformation and subsurface state.

The forecasting windows are defined through a fixed look-back length and a forecast horizon. In the harmonised multi-city framework, the data are

organised using a five-year historical window and a multi-horizon forecast setup. This configuration provides enough temporal context to capture delayed or cumulative deformation response while keeping the forecast problem relevant for planning and risk management.

3.7 Preprocessing strategy and physical consistency

The preprocessing adopted in this study was not purely technical. Each choice was designed to improve the physical interpretability of the final models.

First, groundwater variables were canonicalised so that the predictive stream and the physical stream use compatible but sign-consistent representations of the same underlying process. This avoids ambiguity when linking observed depth-to-water to the hydraulic head used in the reduced physical equations.

Second, thickness information was treated in a censor-aware way. This is particularly important in Zhongshan, where part of the thickness archive is right-censored. A naive use of such measurements would risk introducing systematic distortion into the learning process. By explicitly carrying censor information and constructing an effective thickness field, the framework preserves the stratigraphic signal while acknowledging the limitations of the raw observations.

Third, urban forcing was represented through variables that preserve interpretability. In the early machine-learning stage, this role was played by building concentration. In the later framework it was generalised into a shallow urban-load proxy, which is better suited to cross-city comparison and to the reduced geomechanical logic used by the physics-guided model.

Fourth, autoregressive memory was retained through the inclusion of past cumulative subsidence in the dynamic history. This is important because urban subsidence is path-dependent: the present response is partly conditioned by the system state inherited from previous years. Retaining this information helps the model learn not only static susceptibility, but also

temporal evolution.

3.8 Scientific rationale of the data framework

The data framework developed in this chapter reflects the scientific trajectory of the study. The first stage required a rich and operational predictor set for Nansha in order to test whether machine-learning methods could simulate urban subsidence effectively. The second stage required a stricter and more transferable design, because the objective expanded from single-city simulation to physics-guided forecasting across contrasting urban basins.

For this reason, the present chapter does more than list datasets. It defines the modelling contract of the study. The selected variables represent the main pathways through which urban subsidence is expressed: deformation observed from space, groundwater stress, hydro-climatic forcing, subsurface compressibility, and anthropogenic surface loading. Their harmonisation onto a common annual grid creates a dataset that is rich enough for modern predictive modelling while remaining interpretable in hydrogeological and geomechanical terms.

3.9 Summary

This chapter presented the study areas, data sources, harmonisation strategy, feature construction, and target definition used throughout the study. Nansha and Zhongshan were selected as two contrasting urban subsidence regimes within the Pearl River Delta, allowing the modelling framework to be evaluated beyond a single local setting. The dataset integrates InSAR-derived deformation, groundwater observations, rainfall, borehole-derived thickness and lithology, and urban surface-development information. These heterogeneous sources were harmonised onto a common annual grid and reorganised into static, dynamic, and future-known streams suitable for both machine-learning and physics-guided forecasting.

The resulting data framework establishes the common empirical basis for the chapters that follow. The next chapter builds on this dataset

to present the first methodological stage of the study, namely machine-learning-based land subsidence simulation in Nansha. The subsequent chapter extends the same data foundation toward a physics-guided forecasting framework capable of learning not only deformation trajectories, but also physically interpretable effective structure across multiple urban regimes.

Chapter 4 Machine Learning-Based Subsidence Simulation

4.1 Introduction

The first modelling stage of this study was designed to determine whether heterogeneous urban and environmental observations could be translated into a reliable data-driven simulator of land subsidence in Nansha. The objective was not restricted to pointwise prediction. It also involved three closely related tasks: identifying the dominant controls of deformation, constructing a temporally coherent forecasting workflow from irregularly sampled predictors, and testing whether the trained models could support scenario-based subsidence prevention analysis [29, 24-25, 27].

To address these objectives, two machine-learning models with complementary strengths were adopted. The first was the eXtreme Gradient Boosting Regressor (XGBR), which is particularly effective for nonlinear tabular regression and for ranking the importance of predictors through ensemble tree structure [63]. The second was the Long Short-Term Memory (LSTM) network, which is well suited to sequential learning and to representing delayed temporal dependence in evolving geophysical systems [64-66]. The use of both models makes it possible to compare a strong tree-based learner with a recurrent sequence model rather than relying on a single modelling family.

4.2 Formulation of the prediction problem

Let Ω denote the set of spatial grid cells covering the Nansha study area, and let k denote the weekly time index after temporal interpolation. For each grid cell $i \in \Omega$ and week k , define the predictor vector as

$$\mathbf{x}_{i,k} = [\text{Lith}_i, \text{FSt}_{i,k}, \text{SLd}_{i,k}, \text{BC}_{i,k}, \text{GWL}_{i,k}]^T, \quad (4-1)$$

where Lith is the geological class, FSt is filled-soil thickness, SLd is sand-layer distribution, BC is building concentration, and GWL is groundwater level. These variables were introduced in Chapter 3 as the main descrip-

tors of subsurface condition, anthropogenic loading, and hydrogeological forcing in Nansha [29, 62, 40].

Let $S_{i,k}$ denote cumulative subsidence at grid cell i and week k , derived from the fused InSAR deformation sequence. Instead of directly predicting the cumulative signal, the machine-learning stage models the weekly subsidence increment

$$r_{i,k} = S_{i,k} - S_{i,k-1}. \quad (4-2)$$

This formulation is advantageous because it transforms the forecasting task into the prediction of temporal evolution rather than the direct extrapolation of an accumulated quantity.

The learning problem can then be expressed as the approximation of a nonlinear operator \mathcal{F} such that

$$\hat{r}_{i,k+1} = \mathcal{F}(\mathbf{x}_{i,1:k}, r_{i,1:k}), \quad (4-3)$$

where the interpretation of the historical sequence depends on the chosen model. For XGBR, the sequence is first transformed into a supervised regression table. For LSTM, the sequence is processed directly in its temporal order.

Once the future increments have been predicted, cumulative subsidence is reconstructed recursively as

$$\hat{S}_{i,k+h} = S_{i,k} + \sum_{m=1}^h \hat{r}_{i,k+m}, \quad (4-4)$$

where h denotes the forecast horizon in weeks. This decomposition is important because it separates the dynamic learning problem from the final cumulative-map reconstruction.

4.3 Temporal harmonisation through the delta-rate strategy

A major challenge of the Nansha dataset is that the different predictor layers do not share the same temporal sampling. Some variables are available only for a limited number of survey years, whereas the InSAR deformation archive is substantially denser in time. Direct multivariate forecasting

is therefore not possible without temporal harmonisation. To address this issue, the study introduced a delta-rate strategy, denoted Δr , to interpolate time-varying predictors onto a common weekly axis [29].

For a variable D observed at times t_1 and t_2 , the delta-rate is defined as

$$\Delta r(t_s) = \frac{1}{t_s} \left(\frac{D_{\text{obs}}(t_2) - D_{\text{obs}}(t_1)}{t_2 - t_1} \right), \quad (4-5)$$

where t_s is the target subdivision of the interval, for example in weeks. Intermediate values are then generated recursively by

$$D_{\text{obs}}(t + 1) = D_{\text{obs}}(t) + \Delta r. \quad (4-6)$$

This operation was applied to all time-varying numerical predictors, whereas the geological class Lith remained fixed through time.

The practical consequence of this procedure is that annual or multi-year predictor maps are converted into a denser weekly archive that can be used for time-series learning. In the Nansha application, the predictor stack was rasterised on a $60 \text{ m} \times 60 \text{ m}$ grid and expanded into 395 weekly timestamps covering the period from June 2015 to December 2022 [29]. This weekly representation does not imply that all physical drivers evolve linearly at weekly scale. Its role is to provide a coherent temporal scaffold on which the machine-learning models can be trained and later used for long-horizon simulation.

4.4 Extreme Gradient Boosting Regressor

4.4.1 Theory

XGBR is an additive ensemble of regression trees. For an input vector \mathbf{x}_i , the prediction after K boosting rounds is written as

$$\hat{y}_i = \sum_{k=1}^K f_k(\mathbf{x}_i), \quad f_k \in \mathcal{F}, \quad (4-7)$$

where \mathcal{F} denotes the space of regression trees and \hat{y}_i is the predicted response for sample i . At boosting round t , a new tree is added to correct the

residual structure left by the previous ensemble:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + f_t(\mathbf{x}_i). \quad (4-8)$$

This stage-wise correction is one of the main reasons why gradient boosting performs well in nonlinear environmental modelling [63, 67, 25].

The corresponding objective function is

$$\mathcal{L}(\phi) = \sum_{i=1}^n l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k), \quad (4-9)$$

where $l(y_i, \hat{y}_i)$ is the data loss and $\Omega(f_k)$ penalises tree complexity. In standard boosted-tree notation, the regularisation term may be written as

$$\Omega(f_k) = \gamma T_k + \frac{1}{2} \lambda \sum_{j=1}^{T_k} w_j^2, \quad (4-10)$$

where T_k is the number of leaves in tree k and w_j denotes the leaf weights. This penalty reduces overly complex partitions and helps control overfitting.

4.4.2 Role in the Nansha study

Within the Nansha workflow, XGBR served two purposes. First, it acted as a high-capacity nonlinear regressor for simulating subsidence from the interpolated predictor stack. Second, it provided an interpretable ranking of feature importance through tree-split statistics. This was particularly useful for screening the dominant factors associated with subsidence evolution in Nansha, where groundwater level and building concentration emerged as principal controls in the original analysis [29, 62, 18]. An auxiliary Random Forest benchmark was also introduced in the factor-analysis stage in order to verify whether the main importance structure inferred by XGBR remained stable under another tree-ensemble family [68].

4.5 Long Short-Term Memory network

4.5.1 Sequential memory formulation

The LSTM network is designed to preserve and regulate information over long temporal sequences through an internal cell state and three gates: forget, input, and output [64-65]. Given an input sequence $\{\mathbf{x}_k\}$, the gate equations at time step k are written as

$$\mathbf{f}_k = \sigma(\mathbf{W}_f[\mathbf{h}_{k-1}, \mathbf{x}_k] + \mathbf{b}_f), \quad (4-11)$$

$$\mathbf{i}_k = \sigma(\mathbf{W}_i[\mathbf{h}_{k-1}, \mathbf{x}_k] + \mathbf{b}_i), \quad (4-12)$$

$$\tilde{\mathbf{c}}_k = \tanh(\mathbf{W}_c[\mathbf{h}_{k-1}, \mathbf{x}_k] + \mathbf{b}_c), \quad (4-13)$$

$$\mathbf{c}_k = \mathbf{f}_k \odot \mathbf{c}_{k-1} + \mathbf{i}_k \odot \tilde{\mathbf{c}}_k, \quad (4-14)$$

$$\mathbf{o}_k = \sigma(\mathbf{W}_o[\mathbf{h}_{k-1}, \mathbf{x}_k] + \mathbf{b}_o), \quad (4-15)$$

$$\mathbf{h}_k = \mathbf{o}_k \odot \tanh(\mathbf{c}_k), \quad (4-16)$$

where $\sigma(\cdot)$ is the sigmoid function, \mathbf{h}_k is the hidden state, \mathbf{c}_k is the memory cell, and \odot denotes elementwise multiplication. The predicted response is then obtained from the hidden state by

$$\hat{y}_k = \mathbf{W}_y \mathbf{h}_k + \mathbf{b}_y. \quad (4-17)$$

This sequential structure is particularly suitable for subsidence forecasting because land deformation is path-dependent. The response at a given time is governed not only by the current hydro-environmental conditions, but also by the accumulated deformation history and delayed adjustment of the subsurface system. LSTM therefore provides a natural way to encode temporal memory that is only indirectly observed in the predictor stack.

4.5.2 Role in the Nansha study

In the Nansha application, the LSTM model consumed the weekly predictor sequences and learned the nonlinear temporal dependence between the evolving inputs and the subsidence response [29, 27]. Relative to XGBR, its main advantage is that it preserves sequence order explicitly

rather than through a manually engineered supervised table. Its principal limitation is that it is less transparent than tree-based models and may require careful regularisation when the input archive is partly interpolated rather than fully observed.

4.6 Scenario-based prevention through modified delta rates

An important feature of the machine-learning stage is that it was used not only for passive prediction, but also for prevention-oriented scenario simulation. Once a dominant driver had been identified, its delta-rate could be modified to emulate a reduction strategy. The modified delta-rate was defined as

$$\Delta r_{\text{prev}} = \alpha \Delta r, \quad (4-18)$$

where α is a reduction factor. In the Nansha application, the prevention scenario focused on reducing the influence of groundwater level and building concentration by 80%, motivated by their strong contribution to predicted subsidence [29]. The resulting modified predictor trajectories were then reintroduced into the forecasting workflow.

Because the two model families do not respond to feature attenuation in the same way, model-specific prevention controls were also considered. For LSTM, an L_1 penalty was added to the loss:

$$\mathcal{L}_{\text{new}} = \mathcal{L}_{\text{original}} + \alpha \lambda \sum |w|, \quad (4-19)$$

where λ is the regularisation strength and w denotes network weights. For XGBR, the control was expressed through the maximum tree depth:

$$D_{\text{max}}^{\text{new}} = \alpha D_{\text{max}}^{\text{original}}. \quad (4-20)$$

This formulation does not directly change the predictor itself, but reduces the capacity of the tree ensemble to exploit highly complex partitions tied to the attenuated factor. In methodological terms, the prevention analysis is therefore a conditional re-simulation under modified driver evolution rather than a true mechanistic intervention model.

4.7 Data normalisation and supervised sequence construction

Both XGBR and LSTM require numerical variables to be placed on a comparable scale. The study therefore normalised the input and output variables using min–max scaling:

$$X_n = \frac{X_i - X_{\min}}{X_{\max} - X_{\min}}, \quad (4-21)$$

where X_n is the normalised value and X_i is the original value [69-70]. This transformation limits the effect of differing numerical ranges and improves optimisation stability.

For XGBR, the weekly time-series archive was converted into a supervised regression table in which each sample contains the predictor state and the associated subsidence response at the target step. For LSTM, the same archive was retained as an ordered sequence. The two models therefore operate on the same information content, but under different assumptions about how temporal dependence should be represented.

4.8 Calibration and walk-forward validation

A critical methodological choice in this study was the use of walk-forward validation rather than ordinary random cross-validation. In subsidence forecasting, temporal ordering cannot be ignored: data from later periods should not leak into model training when the evaluation is intended to mimic future prediction. Standard random splitting would therefore lead to optimistic error estimates.

Let the ordered weekly archive be divided into an earlier segment for model fitting and a later segment for out-of-sample validation. The training and validation procedure then advances through time, each time fitting the model on the historical block and evaluating it on the next unseen portion. This strategy preserves causality and provides a more realistic measure of generalisation under temporal drift [29, 69]. In the original Nansha study, the data were partitioned into 70% for training and 30% for validation, with walk-forward evaluation used to avoid the optimistic bias associated with

表 4-1 Hyperparameter configuration used for the machine-learning subsidence simulation in Nansha.

Model	Selected hyperparameters
XGBR	$n_{\text{estimators}} = 1000$; maximum depth = 27; learning rate = 0.1; loss = squared error.
LSTM	number of hidden units = 200; maximum epochs = 100; optimizer = Adamax; dropout = 0.4; learning rate = 0.01.

random folds.

The final hyperparameter configuration adopted in the Nansha study is summarised in Table 4-1. These values were obtained through trial-and-error calibration within the study design reported in Liu et al. [29].

4.9 Evaluation metrics

The predictive performance of the models was assessed using three standard regression metrics: the mean absolute error (MAE), the root-mean-square error (RMSE), and the coefficient of determination R^2 [69, 29]. These are defined as

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \quad (4-22)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad (4-23)$$

and

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}, \quad (4-24)$$

where y_i and \hat{y}_i are the observed and predicted responses, respectively, and \bar{y} is the sample mean of the observations.

In addition to these scalar metrics, the study also used the Taylor diagram as a compact diagnostic of pattern agreement between predictions and observations [71]. For a prediction vector $\hat{\mathbf{y}}$ and observation vector \mathbf{y} , the centred root-mean-square difference is

$$E' = \sqrt{\sigma_y^2 + \sigma_{\hat{y}}^2 - 2\sigma_y\sigma_{\hat{y}}\rho}, \quad (4-25)$$

where σ_y and $\sigma_{\hat{y}}$ are the standard deviations of the observed and predicted series, respectively, and ρ is their correlation coefficient. This representation is useful because it distinguishes a model that captures the overall spatio-temporal pattern from one that merely reduces the average error.

4.10 Methodological significance of the machine-learning stage

The machine-learning stage established three important methodological results.

First, it demonstrated that heterogeneous urban predictors can be assembled into a coherent forecasting dataset despite irregular temporal sampling. The Δr strategy provided a practical solution for bringing geology-informed, hydrogeological, and urban-development variables onto a common temporal support.

Second, it showed that two distinct learning families, namely boosted trees and recurrent neural networks, can both recover meaningful subsidence behaviour from the Nansha archive. This reduces the risk that the main conclusions depend on a single model architecture.

Third, it provided an initial route from explanation to action. By identifying dominant drivers and propagating modified delta-rates through the simulator, the study translated predictive modelling into a prevention-oriented scenario tool. Although this remains an empirical prevention framework rather than a full physical intervention model, it was sufficient to reveal the limits of purely statistical forecasting and to motivate the transition toward a more physically structured framework.

4.11 Limitations of the purely data-driven stage

Despite its strong predictive value, this machine-learning stage leaves several scientific questions unresolved. The first limitation concerns physical interpretability. Although the models can identify influential predictors, they do not explicitly recover hydrogeological fields or compaction timescales. As a result, the forecast may be accurate while remaining only

loosely connected to the underlying mechanics of subsidence.

The second limitation concerns extrapolation under changing forcing. The Δr strategy is effective as a temporal harmonisation tool, but it remains an interpolation-based scaffold. Long-horizon scenarios constructed from it therefore depend on assumptions regarding the persistence of observed trends.

The third limitation concerns uncertainty structure. The machine-learning framework provides deterministic forecasts and performance diagnostics, but it does not impose explicit physical consistency on the predicted trajectories. This makes it difficult to distinguish a statistically acceptable forecast from one that is also geomechanically plausible.

4.12 Summary

This chapter presented the first methodological stage of the study, namely machine-learning-based subsidence simulation in Nansha. Using heterogeneous geological, hydrogeological, and urban-development predictors, the study constructed a temporally harmonised forecasting dataset and evaluated two complementary models, XGBR and LSTM. The results showed that machine-learning methods can recover meaningful subsidence behaviour, identify dominant controls such as groundwater level and building concentration, and support prevention-oriented scenario analysis through modified delta-rate simulation.

At the same time, the chapter clarified the limitations of a purely data-driven approach. Although predictive performance is strong, the forecast remains weakly connected to explicit hydrogeological and geomechanical structure. This limitation motivates the next methodological stage, in which predictive strength is combined with physical accountability through a physics-guided forecasting framework.

Chapter 5 Physics-Guided AI for Subsidence Forecasting

5.1 Introduction

Chapter 4 showed that machine-learning models can recover strong predictive signal from heterogeneous urban and hydro-environmental data. This result is important, but it does not fully resolve the central scientific difficulty of subsidence forecasting. Land subsidence is not merely a pattern-recognition problem. It is the surface manifestation of a coupled hydrogeological and geomechanical system in which groundwater decline, effective stress increase, delayed consolidation, lithological structure, and anthropogenic loading interact over time [12-13, 8, 17]. A forecasting model judged only by point accuracy may therefore remain scientifically incomplete if its internal representation is physically opaque or if its predictions drift toward solutions that are statistically acceptable but mechanically implausible.

This limitation motivates the second modelling stage of the study. The objective is no longer only to predict future cumulative subsidence, but to embed the forecast within a reduced physical scaffold that remains compatible with observed hydrogeological forcing and with interpretable effective subsurface structure. The framework developed for this purpose is GeoPriorSubsNet, a physics-guided forecasting model that couples a probabilistic spatio-temporal predictive backbone to a closure-aware physics head. Its role is to infer not only future subsidence trajectories, but also a set of spatially varying effective fields that provide a mechanistically auditable explanation of the forecast [35, 30-33].

The model is intentionally positioned between two extremes. It is not a purely black-box predictor, because the learned dynamics are constrained through a reduced hydro-geomechanical balance. It is also not a fully resolved numerical basin simulator, because the available data do not justify a highly parameterised inversion at city scale. Instead, the framework uses a reduced but interpretable closure in which effective hydraulic conductivity K , specific storage S_s , and relaxation timescale τ are learned jointly, while

the effective compressible thickness H_{eff} is supplied by the harmonised dataset rather than inferred as a free variable. This design reduces degeneracy, preserves geological interpretability, and makes the forecast physically auditable.

Before introducing the structure of the model, it is useful to summarise the complete methodological pipeline. The workflow shown in Figure 5-1 provides this global view. It links data assembly, harmonization, feature construction, physics-guided learning, calibration, and final forecasting outputs within a single coherent framework. This overview is especially useful because GeoPriorSubsNet is not only a neural forecasting architecture; it is a full modelling system in which the meaning of the forecast depends on how data preparation, physical constraints, and uncertainty evaluation are connected.

As shown in Figure 5-1, the forecasting pipeline begins with heterogeneous data sources, including InSAR subsidence observations, groundwater and rainfall information, borehole lithology and thickness records, and urban surface descriptors such as imperviousness, road density, and low-rise footprints. These inputs are then harmonised through co-registration, annual aggregation, masking, quality control, lithology harmonization, censor-aware thickness construction, and derivation of the shallow urban-load proxy U^* . The harmonised archive is subsequently reorganised into three modelling streams: static features, dynamic features, and future-known drivers. In parallel, the effective thickness field H_{eff} is retained as a dedicated physics input aligned with the forecast horizon.

The core of the workflow is GeoPriorSubsNet itself. The model combines an encoder–decoder forecasting backbone with dual outputs. The data head predicts cumulative subsidence, hydraulic head, and their forecast quantiles, whereas the physics head evaluates reduced consolidation and groundwater-flow consistency. These components are trained jointly through a composite objective that combines quantile-based data loss with physics loss and regularisation terms. After training, the framework supports post-hoc calibration, interval evaluation, physical diagnostics, and

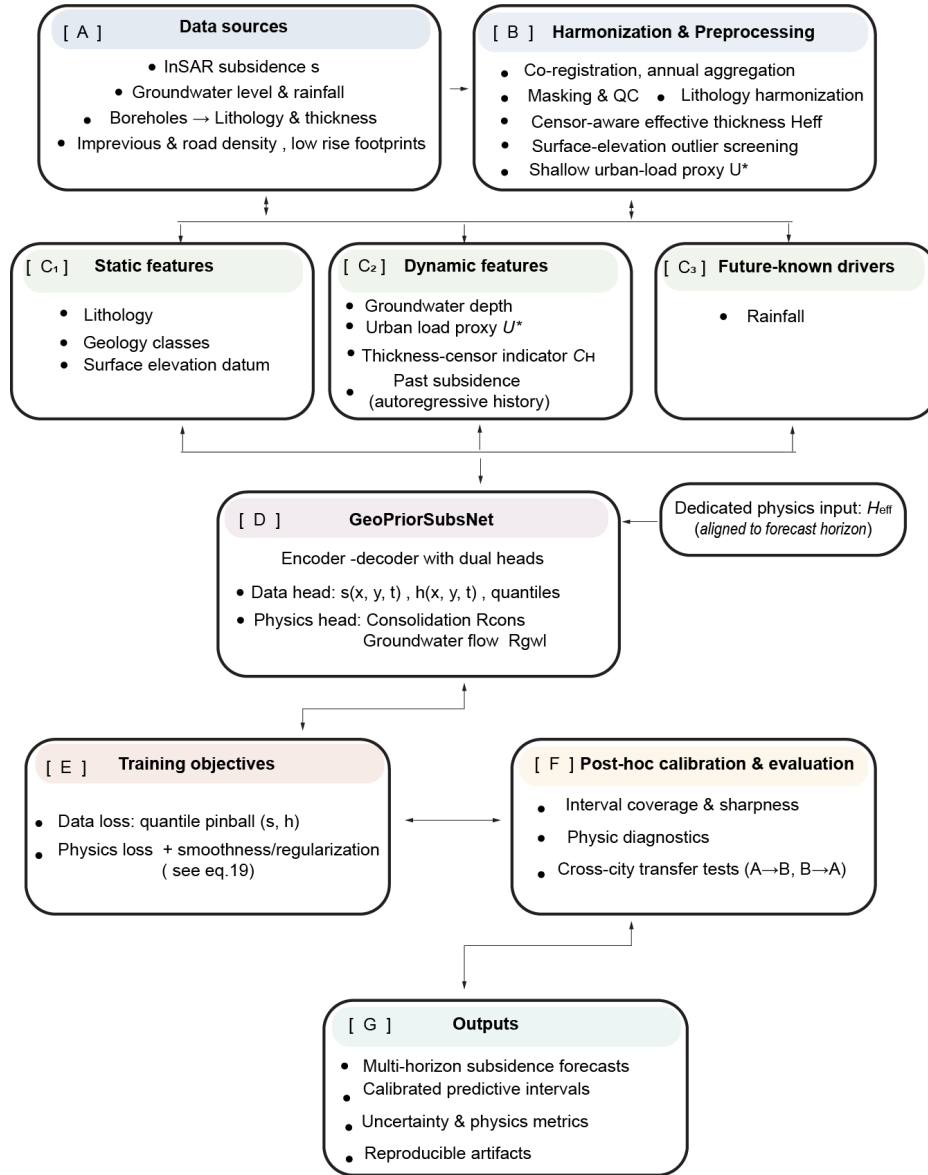


图 5-1 Overall workflow of GeoPriorSubsNet. The framework integrates heterogeneous data sources, harmonization and preprocessing, construction of static and dynamic predictor streams, future-known exogenous forcing, and a dedicated physics input through H_{eff} . These components feed the GeoPriorSubsNet architecture, which combines a probabilistic data head and a physics head within a joint training objective. The workflow further includes post-hoc calibration, physical diagnostics, cross-city transfer tests, and multi-horizon forecasting outputs.

cross-city transfer experiments. The final outputs are therefore not limited to deterministic subsidence maps; they include multi-horizon forecasts, calibrated predictive intervals, uncertainty metrics, physics diagnostics, and reproducible artefacts.

This workflow clarifies an important conceptual point. The predictive component and the physical component are not independent modules

placed side by side. They are coupled through the same latent representation and through a common optimisation problem. Consequently, the value of GeoPriorSubsNet lies not only in its final accuracy, but in the fact that the resulting forecast is physically structured from the moment the data are organised to the moment the outputs are interpreted.

5.2 Forecasting task and input representation

Let $i \in \{1, \dots, |\Omega|\}$ index spatial grid cells and let t denote the forecast origin year. Following Chapter 3, the harmonised archive provides annual cumulative subsidence $s(i, t)$, groundwater depth $d(i, t)$, hydraulic head $h(i, t)$, rainfall $R(i, t)$, shallow urban-load proxy $U^*(i, t)$, surface elevation $z_{\text{surf}}(i)$, and censor-aware effective compressible thickness $H_{\text{eff}}(i)$.

The annual groundwater depth is treated as the driver quantity in the dynamic input stream, with the convention

$$d(i, t) \geq 0, \quad (5-1)$$

positive downward. Hydraulic head is reconstructed from the surface datum as

$$h(i, t) = z_{\text{surf}}(i) - d(i, t), \quad (5-2)$$

so that lower heads correspond to deeper groundwater conditions. The drawdown used in the consolidation formulation is then

$$\Delta h(i, t) = h_{\text{ref}}(i) - h(i, t), \quad (5-3)$$

with $\Delta h > 0$ representing head loss relative to a reference state.

For each grid cell i and forecast origin t , the model forms a look-back window of length L and predicts a forecast horizon of length H . The input tuple is written as

$$X_{i,t} = \left(x_i^{(\text{sta})}, X_{i,t}^{(\text{dyn})}, X_{i,t}^{(\text{fut})}, C_{i,t}, H_{\text{eff}}(i) \right), \quad (5-4)$$

where the static vector is

$$x_i^{(\text{sta})} = [e_{\text{lith}}(i), e_{\text{class}}(i), z_{\text{surf}}(i)], \quad (5-5)$$

the dynamic history is

$$x_{i,t}^{(\text{dyn})} = [d(i, t), s(i, t), R(i, t), U^*(i, t), c_H(i)], \quad (5-6)$$

and the future-known decoder stream is

$$x_{i,t+\tau}^{(\text{fut})} = [R(i, t + \tau)], \quad \tau = 1, \dots, H. \quad (5-7)$$

Here, $c_H(i)$ denotes the censor indicator attached to the thickness field. Future rainfall is supplied as an exogenous decoder input, whereas future groundwater head is not prescribed and must be forecast jointly from the historical sequence and the exogenous forcing.

The supervised targets are the future cumulative-subsidence path and, where available, the future hydraulic-head path,

$$y_i^{(s)} = [s(i, t + 1), \dots, s(i, t + H)], \quad y_i^{(h)} = [h(i, t + 1), \dots, h(i, t + H)]. \quad (5-8)$$

With this input–output contract established, the next step is to define the predictive core of the model. The following subsection therefore introduces the hybrid attentive recurrent backbone, denoted HALNet, which transforms the multi-stream input into horizon-specific latent states and conditional forecast quantiles.

5.3 General architecture of GeoPriorSubsNet

GeoPriorSubsNet is composed of two strongly connected components. The first is a probabilistic predictive backbone that transforms the multi-stream input into horizon-specific latent states and forecast quantiles. The second is a physics head that maps the same latent representation to effective hydrogeological fields and evaluates their consistency through reduced physical residuals. The overall design is illustrated in Figure 5-2.

This structure is designed to preserve three properties simultaneously. First, the model must retain the predictive flexibility of modern sequence learning. Second, it must remain compatible with physically meaningful closure relationships. Third, it must produce uncertainty estimates that are not detached from the physical scaffold of the system. For this reason, the

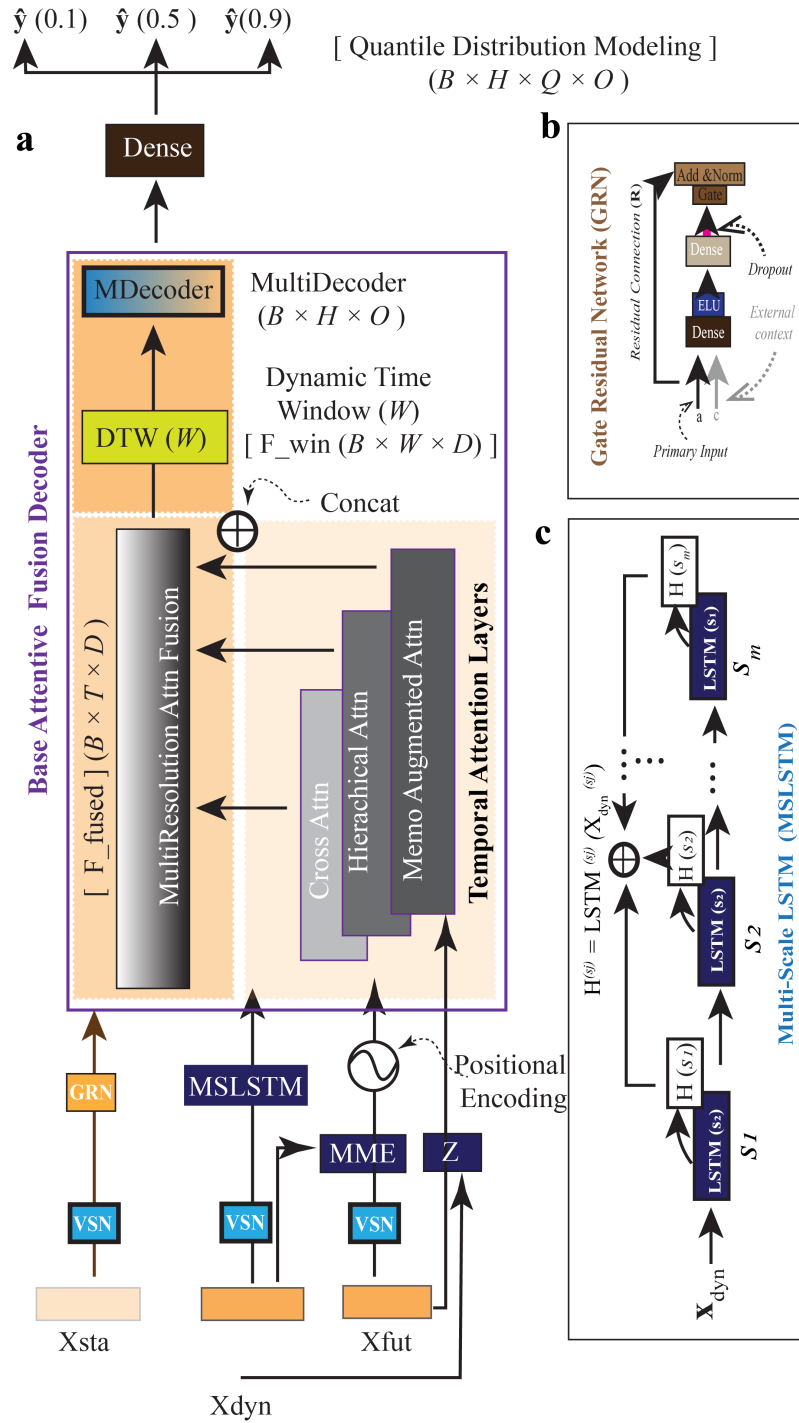


图 5-2 Overview of the GeoPriorSubsNet forecasting architecture. The model combines a hybrid attentive recurrent predictive backbone, denoted HALNet, with a physics head constrained by reduced hydro-geomechanical dynamics. Static, dynamic, and future-known inputs are embedded and fused into horizon-specific latent states. These latent states feed both probabilistic forecast heads for subsidence and hydraulic head, and a physics head that infers effective hydraulic conductivity, specific storage, and closure-consistent relaxation timescale. The framework therefore couples predictive learning, uncertainty quantification, and physical consistency within a single architecture.

data head and the physics head are not trained independently; they are optimised jointly through a composite objective.

5.4 Predictive backbone: HALNet

5.4.1 Probabilistic formulation

The predictive branch of GeoPriorSubsNet is a hybrid attentive recurrent backbone, denoted HALNet, that maps the input tuple Eq. (5-4) to conditional forecast quantiles for subsidence and hydraulic head. Let $\mathcal{Q} = \{q_1, \dots, q_{|\mathcal{Q}|}\}$ be the set of quantile levels, for example $\{0.1, 0.5, 0.9\}$. For each forecast horizon $h \in \{1, \dots, H\}$ and forecast year $t_h = t + h$, the model seeks

$$\hat{s}_q(i, t_h) \approx Q_q[s(i, t_h) | X_{i,t}], \quad \hat{h}_q(i, t_h) \approx Q_q[h(i, t_h) | X_{i,t}]. \quad (5-9)$$

This probabilistic structure is essential because decision-oriented subsidence forecasting depends not only on the expected value of future deformation, but also on threshold exceedance, interval coverage, and horizon-wise sharpness.

5.4.2 Embedding and recurrent encoding

Each input stream is first projected into a common latent dimension d . For the static stream,

$$z_i^{(\text{sta})} = f_{\text{sta}}\left(x_i^{(\text{sta})}\right) \in \mathbb{R}^d, \quad (5-10)$$

for the dynamic stream,

$$z_{i,t}^{(\text{dyn})}(\ell) = f_{\text{dyn}}\left(X_{i,t}^{(\text{dyn})}(\ell), z_i^{(\text{sta})}\right), \quad \ell = 1, \dots, L, \quad (5-11)$$

and for the future-known stream,

$$z_{i,t}^{(\text{fut})}(h) = f_{\text{fut}}\left(X_{i,t}^{(\text{fut})}(h), z_i^{(\text{sta})}\right), \quad h = 1, \dots, H. \quad (5-12)$$

The embedded dynamic history is then processed by an LSTM-based recurrent encoder,

$$u_{i,t}(\ell) = \text{RNN}\left(z_{i,t}^{(\text{dyn})}(\ell), u_{i,t}(\ell - 1)\right), \quad \ell = 1, \dots, L. \quad (5-13)$$

This recurrent representation preserves temporal ordering and allows the forecast to depend on delayed system memory, which is essential for subsidence because present deformation depends on previously accumulated compaction and past groundwater evolution [65-66, 35].

5.4.3 Attentive fusion and probabilistic decoder

For each forecast step h , the recurrent history, static context, future-known forcing, and forecast-aligned coordinates are fused into a horizon-specific latent state

$$f_{i,t}(h) = \text{AttnFuse}\left(\{u_{i,t}(\ell)\}_{\ell=1}^L, z_i^{(\text{sta})}, z_{i,t}^{(\text{fut})}(h), C_{i,t}\right) \in \mathbb{R}^{d_f}, \quad (5-14)$$

where $C_{i,t}$ contains the forecast-aligned space–time coordinates (t_h, x_i, y_i) for each horizon step.

The fused state is mapped through a lightweight decoder,

$$d_{i,t}(h) = g_{\text{dec}}(f_{i,t}(h)) \in \mathbb{R}^{d_d}, \quad (5-15)$$

and the quantile outputs are then given by affine heads,

$$\hat{s}_q(i, t_h) = \mathbf{w}_q^{(s)\top} d_{i,t}(h) + b_q^{(s)}, \quad \hat{h}_q(i, t_h) = \mathbf{w}_q^{(h)\top} d_{i,t}(h) + b_q^{(h)}. \quad (5-16)$$

The predictive complexity is therefore concentrated in the attentive fusion backbone, while the output heads remain intentionally simple so that they can be coupled cleanly with the physics penalties.

5.5 Physics head and effective closure fields

5.5.1 Need for a dedicated physics head

A probabilistic data head can learn conditional forecast distributions, but it cannot ensure that the learned trajectories remain compatible with reduced hydro–geomechanical dynamics. GeoPriorSubsNet therefore attaches a dedicated physics head to the same latent state $f_{i,t}(h)$ used by the decoder. This physics head does not predict subsidence directly. Instead, it maps the latent representation to effective closure fields capable of explaining the forecast through reduced physical laws.

The raw physics output is

$$p_{i,t}(h) = \left[\tilde{K}_{i,t}(h), \tilde{S}_{s,i,t}(h), \widetilde{\delta \log \tau}_{i,t}(h), \tilde{Q}_{i,t}(h) \right], \quad (5-17)$$

corresponding respectively to hydraulic conductivity, specific storage, a log-timescale offset, and an optional effective source/sink term. A key modelling decision is that the thickness field is *not* learned by the physics head. Instead, $H_{\text{eff}}(i)$ is read directly from the harmonised data and used as a dedicated geological input. This removes an unnecessary degree of freedom and keeps thickness interpretable as an externally constrained stratigraphic control.

5.5.2 Drainage thickness and prior-centred fields

The effective drainage thickness is defined from the supplied thickness field as

$$H_d(i) = \alpha_H H_{\text{eff}}(i), \quad (5-18)$$

where $\alpha_H \in (0, 1]$ is a drainage-thickness factor. The learned conductivity and storage fields are then expressed in log space around lithology-informed prior baselines:

$$\log K(i, t_h) = \log K_{\text{prior}}(i) + \delta K(i, t_h), \quad (5-19)$$

$$\log S_s(i, t_h) = \log S_{s,\text{prior}}(i) + \delta S_s(i, t_h). \quad (5-20)$$

These priors do not represent laboratory truth. Their role is to anchor the inversion within geologically admissible ranges and prevent the learned fields from drifting arbitrarily.

5.5.3 Closure-consistent relaxation timescale

The closure-implied relaxation timescale is computed from H_d , K , and S_s as

$$\tau_{\text{cl}}(i, t_h) = \frac{H_d(i)^2 S_s(i, t_h)}{\pi^2 \kappa_b K(i, t_h)}, \quad (5-21)$$

or equivalently, in log form,

$$\log \tau_{\text{cl}}(i, t_h) = 2 \log H_d(i) + \log S_s(i, t_h) - \log(\pi^2) - \log \kappa_b - \log K(i, t_h). \quad (5-22)$$

Here, $\kappa_b > 0$ is a dimensionless consistency factor.

The model does not force the learned relaxation time to coincide exactly with τ_{cl} . Instead, it learns an additive discrepancy in log space,

$$\delta \log \tau(i, t_h) = \widetilde{\delta \log \tau}_{i,t}(h), \quad (5-23)$$

so that

$$\log \tau(i, t_h) = \log \tau_{\text{cl}}(i, t_h) + \delta \log \tau(i, t_h). \quad (5-24)$$

This is a central component of the framework. The closure relation provides the physical scaffold, while $\delta \log \tau$ absorbs structured model discrepancy, unresolved heterogeneity, and departures from the idealised one-dimensional relaxation picture. The timescale is therefore closure-anchored rather than completely rigid or completely free.

5.5.4 Positivity and admissible bounds

Because K , S_s , and τ must remain positive, the guarded log fields are mapped through an exponential transform,

$$q(i, t_h) = \exp(\ell_q^*(i, t_h)) + \varepsilon_q, \quad q \in \{K, S_s, \tau\}, \quad (5-25)$$

where $\varepsilon_q > 0$ is a small numerical guard. In hard-bounds mode, the guarded log-value is clipped to an admissible interval,

$$\ell_q^* = \min(\max(\ell_q, \ell_{q,\min}), \ell_{q,\max}). \quad (5-26)$$

This bounded parameterisation is necessary because positivity alone would still allow implausible trade-offs, such as very low conductivity being compensated by extremely high storage or relaxation time.

5.6 Reduced hydro–geomechanical dynamics

5.6.1 Groundwater residual

The first physics component is a reduced groundwater-flow residual, written as

$$R_{\text{gw}}(i, t_h) = S_s(i, t_h) \partial_t h(i, t_h) - \nabla \cdot (K(i, t_h) \nabla h(i, t_h)) - Q(i, t_h), \quad (5-27)$$

where $Q(i, t_h)$ is the optional effective source/sink term predicted by the fourth physics-head channel. This residual does not attempt to resolve a fully detailed basin flow model. Its role is to enforce consistency with an effective diffusion-like representation of head evolution at grid scale.

5.6.2 Consolidation residual

The second, and more central, physical component is the reduced consolidation residual

$$R_{\text{cons}}(i, t_h) = \partial_t s(i, t_h) - \frac{s_{\text{eq}}(i, t_h) - s(i, t_h)}{\tau(i, t_h)}, \quad (5-28)$$

where the equilibrium settlement is approximated by

$$s_{\text{eq}}(i, t_h) \approx S_s(i, t_h) \Delta h(i, t_h) H_{\text{eff}}(i). \quad (5-29)$$

Equation (5-28) states that the instantaneous settlement rate is proportional to the distance between the current state and the equilibrium settlement induced by drawdown, scaled by the local relaxation time τ . This gives the model a transparent physical meaning: the settlement trajectory is interpreted as a delayed relaxation toward a drawdown-dependent equilibrium state.

The use of H_{eff} in Eq. (5-29) is deliberate. It ensures that the data-derived geological thickness enters the equilibrium compaction pathway directly, rather than being absorbed into a less interpretable free parameter. This is a major reason why the framework remains more geologically interpretable than a generic physics-regularised network.

5.6.3 Residual losses

The scalar groundwater and consolidation residual losses are defined as

$$L_{\text{gw}} = \langle \tilde{R}_{\text{gw}}^2 \rangle, \quad L_{\text{cons}} = \langle \tilde{R}_{\text{cons}}^2 \rangle, \quad (5-30)$$

where $\langle \cdot \rangle$ denotes averaging over samples and forecast horizons, and the tildes indicate optionally non-dimensionalised residuals. This non-dimensionalisation is useful because it prevents one family of residuals from dominating the optimisation only because of unit scale.

5.7 Probabilistic data loss and composite objective

5.7.1 Quantile loss

The probabilistic data term is defined through the pinball loss. For a residual e and quantile level q ,

$$\rho_q(e) = (q - \mathbb{I}\{e < 0\}) e, \quad (5-31)$$

where $\mathbb{I}\{\cdot\}$ is the indicator function. For a mini-batch of size B , the subsidence quantile loss is

$$L_s = \frac{1}{B} \sum_{i=1}^B \sum_{h=1}^H \sum_{q \in \mathcal{Q}} w_s(q) \rho_q(s(i, t_h) - \hat{s}_q(i, t_h)), \quad (5-32)$$

and when hydraulic-head supervision is available, the head quantile loss is

$$L_h = \frac{1}{B} \sum_{i=1}^B \sum_{h=1}^H m_{i,t_h} \sum_{q \in \mathcal{Q}} w_h(q) \rho_q(h(i, t_h) - \hat{h}_q(i, t_h)), \quad (5-33)$$

where $m_{i,t_h} \in \{0, 1\}$ is an availability mask. The total data loss is then

$$L_{\text{data}} = L_s + L_h. \quad (5-34)$$

5.7.2 Timescale prior, smoothness, and composite objective

The central regulariser of the model is not a direct penalty on K , S_s , or H_d individually, but a penalty on the mismatch between the learned relaxation time and the closure-implied timescale:

$$R_{\text{prior}}(i, t_h) = \log \tau(i, t_h) - \log \tau_{\text{cl}}(i, t_h), \quad (5-35)$$

with associated loss

$$L_{\text{prior}} = \left\langle R_{\text{prior}}^2 \right\rangle. \quad (5-36)$$

Working in log space is essential because characteristic timescales may vary across orders of magnitude.

To suppress unstable pixel-to-pixel oscillation in the inferred hydraulic fields, a smoothness regularisation is introduced:

$$L_{\text{smooth}} = \left\langle \|\nabla K\|_2^2 + \|\nabla S_s\|_2^2 \right\rangle. \quad (5-37)$$

In addition, bounds penalties are imposed so that the inferred effective parameters remain within physically admissible intervals.

The full training objective is therefore

$$L_{\text{total}} = L_{\text{data}} + \lambda_{\text{cons}}L_{\text{cons}} + \lambda_{\text{gw}}L_{\text{gw}} + \lambda_{\text{prior}}L_{\text{prior}} + \lambda_{\text{smooth}}L_{\text{smooth}} + \lambda_{\text{bounds}}L_{\text{bounds}} + \lambda_Q L_Q, \quad (5-38)$$

where L_Q is an optional regularisation on the effective source term when that channel is active. This objective combines predictive accuracy, mechanical plausibility, closure consistency, and numerical stability in a single optimisation problem.

For model monitoring and selection, two diagnostic root-mean-square errors are especially important:

$$\varepsilon_{\text{cons}} = \sqrt{\langle \tilde{R}_{\text{cons}}^2 \rangle}, \quad \varepsilon_{\text{prior}} = \sqrt{\langle R_{\text{prior}}^2 \rangle}. \quad (5-39)$$

The first measures mismatch to the reduced relaxation law, whereas the second measures mismatch to the closure-implied timescale.

5.8 Identifiability and interpretation of the inferred fields

A critical methodological issue is that the closure Eq. (5-21) contains an intrinsic multiplicative trade-off between conductivity, storage, and drainage thickness. Taking differentials in log space gives

$$\delta \log \tau_{\text{cl}} = 2 \delta \log H_d + \delta \log S_s - \delta \log K - \delta \log \kappa_b. \quad (5-40)$$

If κ_b is fixed, then perturbations satisfying

$$\delta \log K \approx \delta \log S_s + 2 \delta \log H_d \quad (5-41)$$

leave the implied relaxation time nearly unchanged. This means that different combinations of (K, S_s, H_d) may generate very similar effective consolidation behaviour.

For this reason, the framework does not claim unique pixel-scale recovery of all closure components from subsidence alone. The principal recoverable quantity is the effective relaxation timescale τ , whereas the decomposition into K , S_s , and H_d must be interpreted as a prior-regularised

and physically constrained factorisation. The prior and bounds penalties do not weaken the inversion; rather, they make it scientifically stable by restricting the solution to an admissible manifold.

5.9 Methodological significance of the physics-guided stage

The transition from Chapter 4 to the present framework is not simply a change of architecture. It changes the meaning of the forecast itself. In the machine-learning stage, the model predicts subsidence from observed predictors and identifies dominant controls, but it does not explain the forecast through explicit effective hydro–geomechanical structure. In the present stage, the same forecasting problem is reformulated so that the predicted trajectories must remain compatible with a reduced consolidation law, a groundwater-flow residual, and a closure-consistent timescale.

This change has three major consequences. First, it yields interpretable effective fields rather than only statistical importance measures. Second, it produces probabilistic forecasts that can be assessed not only by coverage and sharpness, but also by physics-audit diagnostics. Third, it allows differences between Nansha and Zhongshan to be examined in terms of effective dynamics rather than only data distribution. In this sense, the physics-guided stage does not replace the machine-learning stage; it builds upon it and resolves its principal conceptual limitation.

5.10 Summary

This chapter presented the methodological core of GeoPriorSubsNet, the physics-guided forecasting framework developed in this study. The model combines a probabilistic spatio-temporal predictive backbone, HAL-Net, with a dedicated physics head that infers closure-consistent effective fields and evaluates them through reduced hydro–geomechanical residuals. The formulation explicitly couples forecast quantiles, effective conductivity, effective storage, effective drainage thickness, and relaxation timescale within a single optimisation problem.

The chapter also showed why this framework constitutes the princi-

pal novelty of the study. It goes beyond purely data-driven forecasting by making the predicted trajectories physically auditable and by linking them to interpretable effective structure. At the same time, it avoids the impracticality of a fully resolved basin-scale inversion under limited urban data. The next chapter evaluates this framework experimentally through deterministic accuracy, probabilistic calibration, inferred effective fields, physical diagnostics, external validation, and cross-city transfer.

Chapter 6 Results and Experimental Validation

6.1 Introduction

This chapter evaluates the forecasting framework from six complementary perspectives. First, it examines whether Nansha and Zhongshan occupy similar driver spaces or whether they represent genuinely distinct subsidence regimes. Second, it quantifies deterministic forecasting skill and compares the physics-guided model against strong non-physics and physics-based baselines. Third, it assesses probabilistic calibration, interval sharpness, and horizon stability. Fourth, it examines the inferred effective fields and the extent to which the forecast trajectories remain consistent with the reduced physical scaffold introduced in Chapter 5. Fifth, it tests whether the inferred fields remain anchored to independent local information through point-support validation in Zhongshan. Finally, it evaluates cross-city transfer in order to determine whether the learned representation can be reused under realistic distribution shift.

The central question is not whether the model predicts well in a narrow statistical sense, but whether it produces forecasts that are simultaneously accurate, calibrated, physically auditable, and useful for deployment in contrasting urban basins. This distinction is important for land subsidence, where a modest gain in point accuracy is not sufficient if it is obtained at the expense of uncertainty quality or mechanical plausibility [8, 17, 71].

6.2 Contrasting driver regimes across Nansha and Zhongshan

Before evaluating the forecasts themselves, it is necessary to verify whether Nansha and Zhongshan occupy comparable input spaces or whether they represent distinct urban subsidence regimes. Field observations already indicate that the underlying deformation is severe in both cities, with vertical offsets of approximately 400 mm in Nansha and more than 680 mm in Zhongshan, as documented earlier in Figure 3-1. This confirms that both basins are relevant test beds for risk-oriented subsidence

forecasting.

The harmonised driver analysis shows that the two cities differ not only in deformation magnitude, but also in the distribution of the principal physical and anthropogenic controls. Nansha samples deeper groundwater conditions, a broader rainfall span, and a wider range of the urban-load proxy, whereas Zhongshan is concentrated at shallower head conditions, annual rainfall between roughly 1000 and 3000 mm, and generally higher baseline urban-load values. Lithological composition is also significantly different, with $\chi^2 \approx 4.23 \times 10^4$ and Cramér's $V \approx 0.49$, indicating a strong between-city contrast in the subsurface material mix. In broad terms, Zhongshan contains a larger proportion of more competent units, whereas Nansha contains a higher fraction of fine-grained and more compressible soils. Such differences are consistent with broader observations from deltaic and rapidly reclaimed coastal settings, where compressible fills and shallow sedimentary heterogeneity amplify subsidence sensitivity [16-17, 10].

The empirical driver–response analysis further clarifies these contrasts. In Nansha, cumulative subsidence increases monotonically with depth to water and rises steeply up to approximately 50–60 m before flattening. In Zhongshan, the groundwater-depth range is much shallower and the response weakens beyond about 10–15 m. Rainfall acts as a weak nonlinear mitigating factor in both cities, with its influence more visible in Nansha because the hydro-climatic range is broader there. Thickness effects are positive and saturating in both settings, although the Zhongshan signal contains visible plateaus linked to the known censoring structure of the thickness archive. The shallow urban-load proxy exhibits a lower-slope convex response and mainly modulates the baseline when combined with deep groundwater or thicker compressible sections. These relationships are summarised in Figure 6-1.

These results are important because they show that the two study areas are not redundant replications of the same deformation problem. They represent two distinct regimes of practical significance: a highly heteroge-

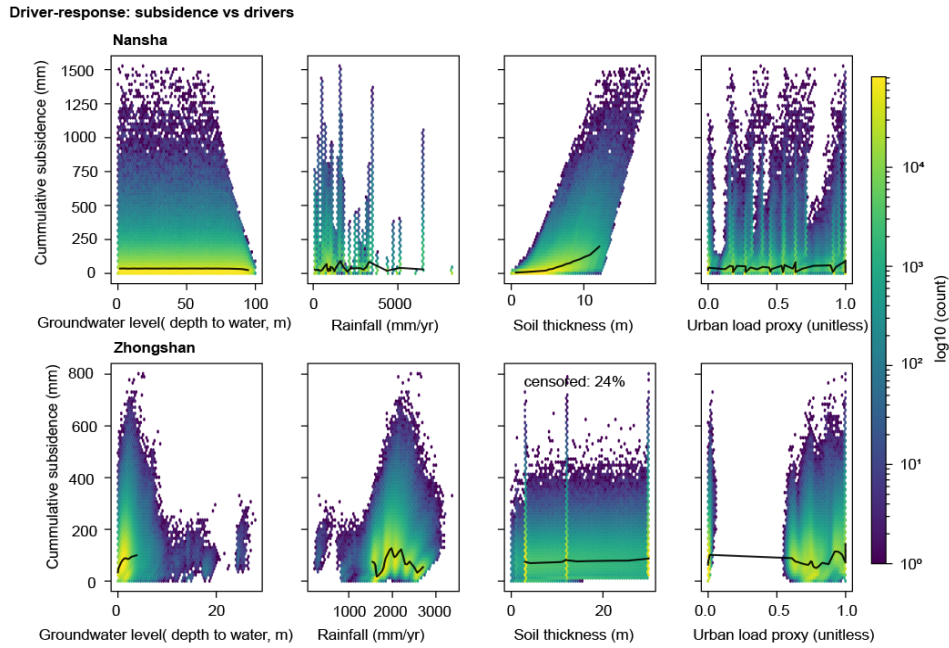


图 6-1 Driver–response relationships between cumulative subsidence and the main controlling variables in Nansha and Zhongshan. The panels show the empirical relation between cumulative subsidence and groundwater depth, rainfall, effective soil thickness, and shallow urban-load proxy. The black curves indicate the smoothed trend. Nansha exhibits broader hydro-environmental variability and stronger monotonic sensitivity to groundwater depth and thickness, whereas Zhongshan shows a shallower groundwater regime and a more compact driver space.

neous, reclaimed coastal system with broad hydro-environmental variability, and a more structured urban basin with clearer hydrostratigraphic organisation. This contrast gives the later benchmarking and transfer results a stronger scientific meaning than a single-city evaluation would allow.

6.3 Deterministic predictive performance

6.3.1 Core out-of-sample accuracy

Across both cities, the full GeoPriorSubsNet configuration achieves strong out-of-sample agreement with observed cumulative subsidence. In Nansha, the model reaches $R^2 = 0.88$ with $\text{MAE} = 9.27$ mm and $\text{MSE} = 262.47$ mm². In Zhongshan, the corresponding performance is $R^2 = 0.90$, $\text{MAE} = 1.78$ mm, and $\text{MSE} = 7.57$ mm². The large difference in absolute error between the two cities is scientifically meaningful. It indicates that the coastal Nansha system is the more difficult forecasting environment, which

is consistent with its broader covariate range, stronger non-stationarity, and greater fine-scale heterogeneity.

A controlled ablation clarifies the role of the physics regularisers. In Nansha, removing the physics terms yields a small improvement in average point error, with MAE decreasing from 9.27 to 9.03 mm and MSE decreasing from 262.47 to 237.78 mm². This shows that in a highly heterogeneous basin, a purely data-driven pathway can still match, and on this split marginally improve, the mean deterministic error. The gain, however, is limited and must be interpreted together with the loss of physical structure. In Zhongshan, the pattern is reversed. Removing physics increases the MAE from 1.78 to 2.14 mm and the MSE from 7.57 to 16.36 mm², showing that in the more organised basin the physics terms act as an effective regulariser that improves generalisation by constraining the temporal response. These contrasts are summarised in Figure 6-2.

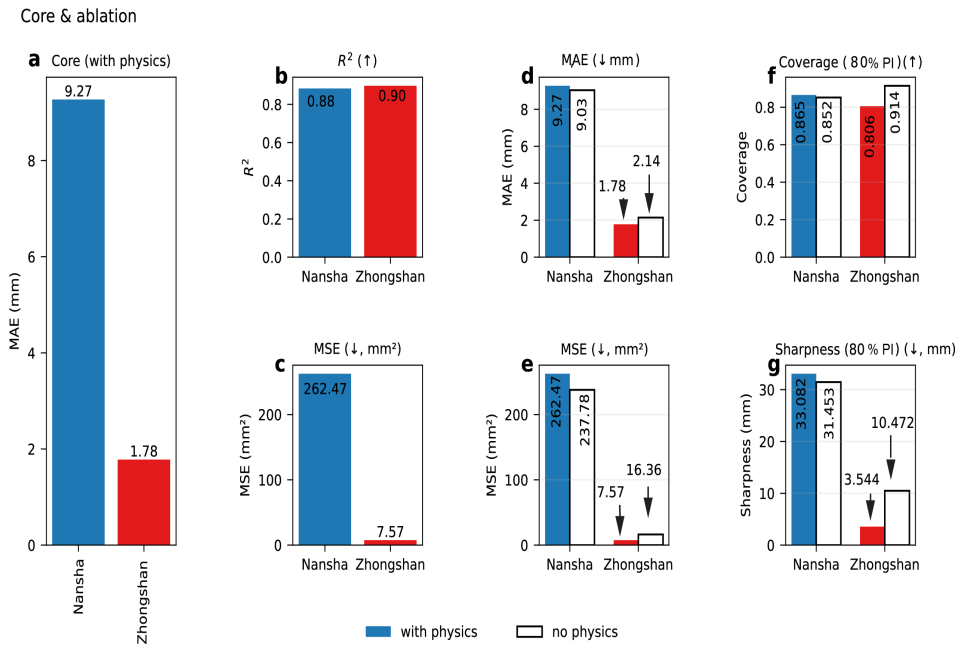


图 6-2 Core deterministic performance and no-physics ablation. The figure summarises the out-of-sample performance of GeoPriorSubsNet in Nansha and Zhongshan, together with the effect of removing the physics constraints. The physics-guided configuration retains strong deterministic skill in both cities and yields a particularly clear advantage in Zhongshan, where it improves both accuracy and uncertainty sharpness relative to the no-physics alternative.

This contrast between the two basins is consistent with the idea that physics guidance is most directly beneficial where the subsurface response

is sufficiently coherent for the reduced closure to act as an informative constraint. In more heterogeneous coastal settings, the data pathway can recover much of the short-horizon correlation structure, but the physical scaffold remains important for interpretability and robustness [21-22, 58].

6.3.2 Benchmark comparison against non-physics and physics-based baselines

The benchmark comparison broadens the interpretation of these results. GeoPriorSubsNet was evaluated against a strong attentive backbone (HALNet), an LSTM sequence model, the Temporal Fusion Transformer (TFT), and a poroelastic PINN-style surrogate (PoroElasticSubsNet). In Nansha, GeoPriorSubsNet attains the lowest deterministic error among all tested models, with MAE = 9.27 mm compared with 11.92 mm for PoroElasticSubsNet, 18.03 mm for HALNet, 19.10 mm for TFT, and 20.21 mm for LSTM. In Zhongshan, the same ordering remains, with GeoPriorSubsNet reaching MAE = 1.78 mm, ahead of PoroElasticSubsNet at 2.74 mm, HALNet at 9.54 mm, TFT at 22.02 mm, and LSTM at 34.87 mm.

The R^2 pattern is similar. In Nansha, GeoPriorSubsNet reaches $R^2 = 0.883$, slightly exceeding the poroelastic surrogate at 0.878 and clearly outperforming the non-physics baselines. In Zhongshan, GeoPriorSubsNet reaches $R^2 = 0.897$, again slightly above the poroelastic surrogate at 0.895 and well above the other baselines. This is notable because it shows that the gain from the physics-guided architecture is not merely relative to purely statistical models; it also remains visible against a more explicitly mechanics-informed surrogate. The comparison suggests that a reduced closure-consistent scaffold may be better suited to city-scale forecasting than a more direct poroelastic PINN formulation when boundary conditions, multi-depth anthropogenic forcing, and basin-wide heterogeneity are only partially observed [30-31, 33].

6.4 Probabilistic calibration, sharpness, and horizon stability

Deterministic accuracy alone is not sufficient for decision-oriented subsidence forecasting. Urban planning and risk management depend on

the quality of the predictive distribution, especially when the relevant decision variable is exceedance probability rather than mean error. For that reason, the uncertainty evaluation focused on central prediction interval coverage, sharpness, reliability at selected quantiles, and horizon-wise stability.

In Nansha, the nominal 80% prediction interval produced by GeoPriorSubsNet attains empirical coverage of 0.865 with sharpness 33.08 mm. In Zhongshan, the interval coverage is 0.806, much closer to nominal, with substantially sharper intervals of width 3.54 mm. These results show that the uncertainty structure adapts to basin complexity. Nansha requires broader intervals because the deformation regime is more variable and more weakly constrained. Zhongshan supports much sharper intervals while remaining close to the nominal target.

The comparison with the no-physics ablation is especially revealing. In Nansha, the no-physics model achieves similar coverage (0.852) with slightly tighter intervals (31.45 mm), indicating that nominal calibration can be matched in this basin even without explicit physical constraints. In Zhongshan, however, the difference is much larger: the no-physics variant over-covers at 0.914 but does so with intervals nearly three times wider (10.47 mm) than those of the physics-guided model. This means that the non-physics model appears conservative, but only by diffusing predictive mass across a broad range of outcomes. The physics-guided model is preferable because it produces intervals that are both more informative and more closely aligned with the target coverage.

The reliability diagrams provide a more detailed view of this behaviour. In Nansha, the lower quantile is mildly conservative, whereas the upper quantile remains close to nominal. In Zhongshan, the main deviation appears near the median, where the central uncertainty is somewhat under-dispersed. Even so, the upper quantile again remains closer to nominal than the central part of the distribution. This is practically important because many policy questions depend more strongly on upper-tail exceedance than on central prediction. These features are illustrated in Figure 6-3.

A further result of practical importance is that calibration remains sta-

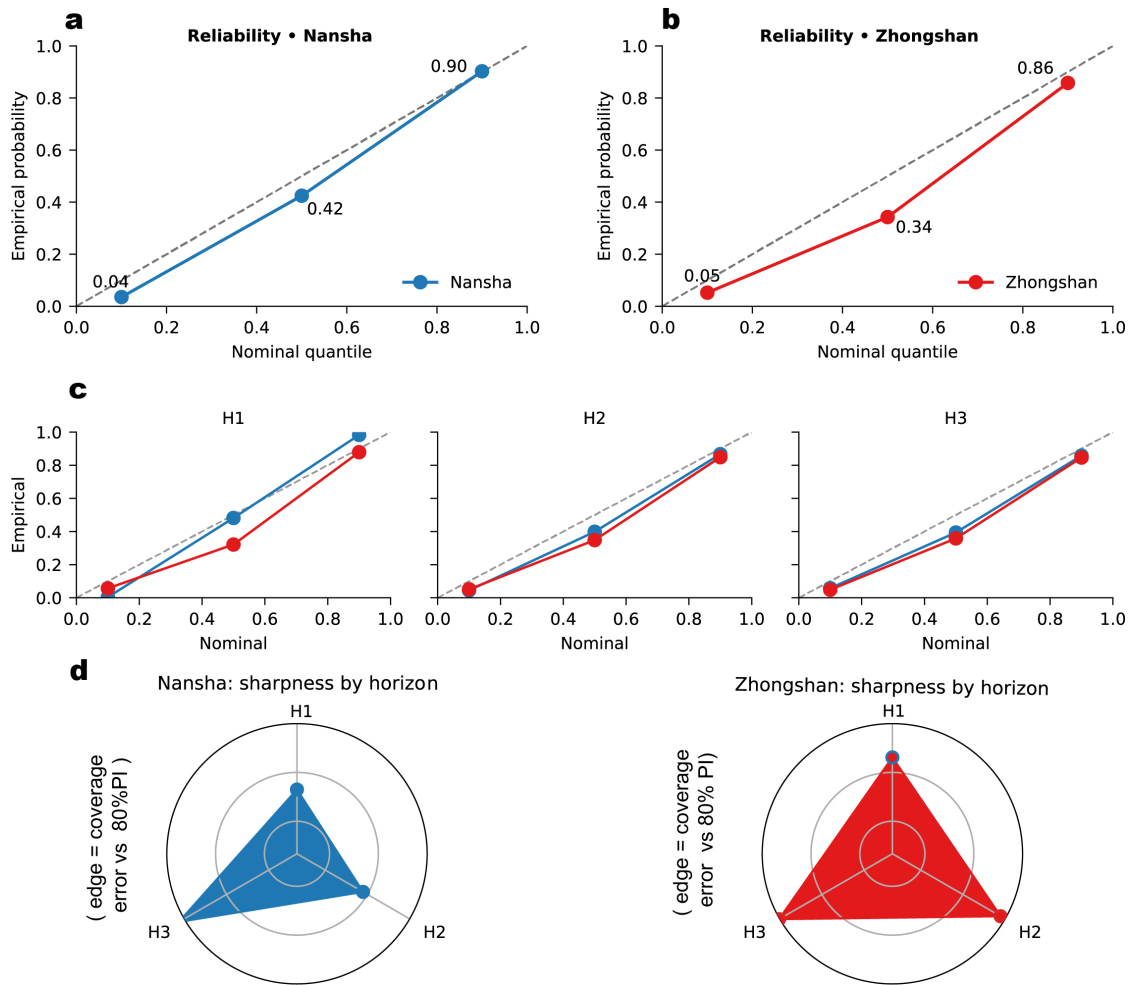


图 6-3 Reliability, horizon-wise calibration, and sharpness. Panels a–b show the reliability curves for Nansha and Zhongshan. Panel c shows the horizon-specific reliability behaviour from H1 to H3. Panel d summarises the horizon-wise sharpness. The results indicate that the predictive distribution remains reasonably well calibrated across horizons, with broader intervals in Nansha and sharper but still well-targeted intervals in Zhongshan.

ble across forecast horizons. The horizon-specific reliability curves show similar qualitative behaviour from H1 to H3, with only modest departures from the ideal diagonal. The model therefore avoids a common failure mode in multi-step forecasting, where short-horizon intervals appear calibrated but long-horizon intervals become sharply overconfident.

This conclusion is reinforced by the PIT histograms, the horizon-wise coverage profiles, the sharpness curves, and the quantile-residual summaries shown in Figure 6-4. The 80% interval coverage remains close to target across lead times, whereas the increase in interval width is city-dependent, with stronger uncertainty growth in the more heterogeneous

coastal setting.

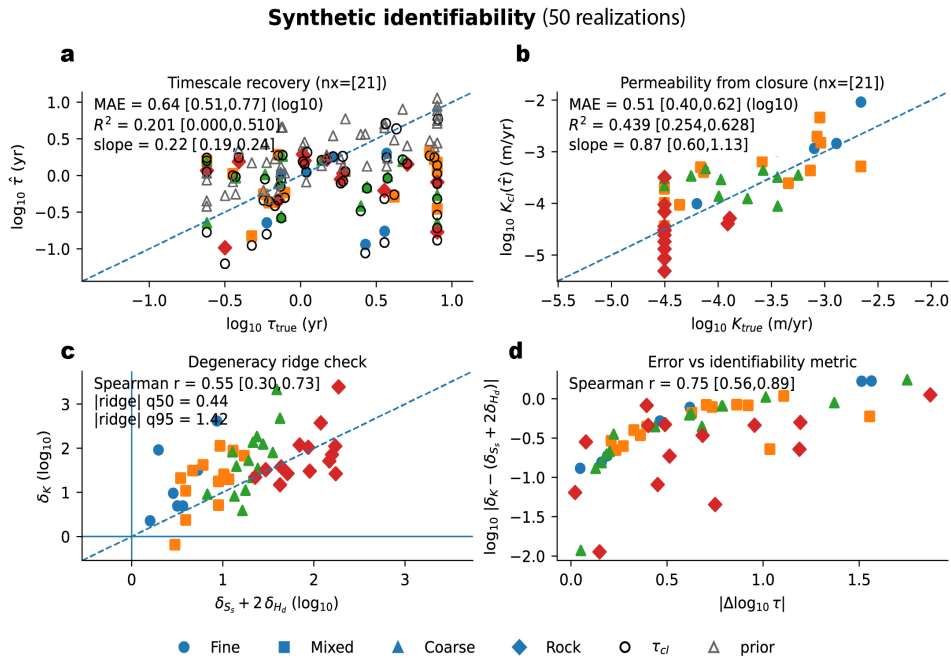


图 6-4 Additional probabilistic diagnostics. The figure reports PIT histograms for Nansha and Zhongshan, horizon-wise coverage against the nominal 80% target, sharpness evolution with lead time, calibration factors per horizon, and quantile residuals across horizons. Together these diagnostics confirm that uncertainty behaviour remains stable with lead time, although Nansha requires broader intervals than Zhongshan.

6.5 Forecasted hotspot evolution and persistence

Beyond calibration at the point level, decision-making also depends on the spatial organisation of forecast risk. The horizon-wise hotspot maps show that Nansha contains broader and more persistent high-risk zones, whereas Zhongshan displays more localised and less persistent hotspot structures. This difference is visible both in the maps of projected annual change and in the exceedance-probability fields for $|s| \geq 50 \text{ mm yr}^{-1}$. The cluster-level summaries further show that Nansha contains several high-priority clusters with large persistence scores, while Zhongshan is dominated by one leading cluster followed by a steeper drop in mean risk. These forecasted hotspot structures are shown in Figure 6-5.

This risk-oriented view complements the scalar accuracy metrics. It shows that the model is not only accurate on average, but also able to organise the forecast into spatially meaningful priority structures, which is

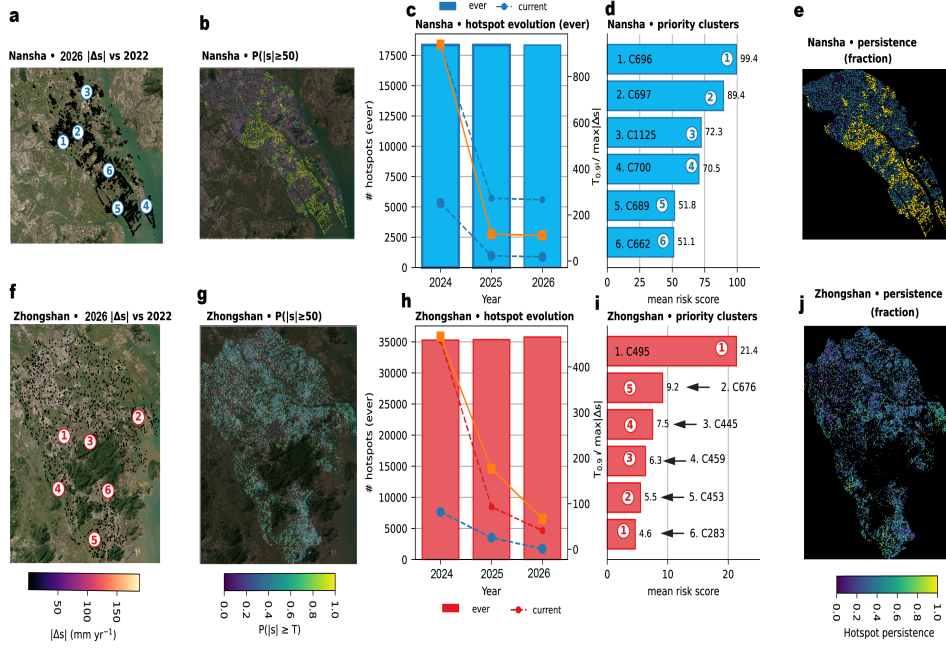


图 6-5 Forecasted hotspot evolution and persistence in Nansha and Zhongshan. Panels a–b and f–g show the forecasted annual subsidence change and exceedance probability maps. Panels c–d and h–i summarise hotspot counts, temporal evolution, and priority clusters. Panels e and j show hotspot persistence. The results indicate broader and more persistent risk corridors in Nansha and more localised hotspot organisation in Zhongshan.

essential for subsidence monitoring and intervention planning.

6.6 Inferred effective fields and physics diagnostics

6.6.1 Closure consistency and dynamic balance

A central scientific claim of the physics-guided framework is that the forecast is not merely regularised by physics, but remains auditable against explicit dynamic and closure diagnostics. The first of these is the consolidation residual R_{cons} , which measures the mismatch between the predicted settlement rate and the reduced relaxation law introduced in Chapter 5. In both cities, the scaled residual remains sharply centred on zero. The corresponding dimensionless RMS diagnostic is $\varepsilon_{\text{cons}} = 2.61 \times 10^{-6}$ in Nansha and 9.56×10^{-7} in Zhongshan. These values show that the learned subsidence trajectories remain tightly anchored to the rate-form consolidation balance. The distribution of R_{cons} is shown in Figure 6-6.

The second diagnostic concerns the timescale closure. The learned re-

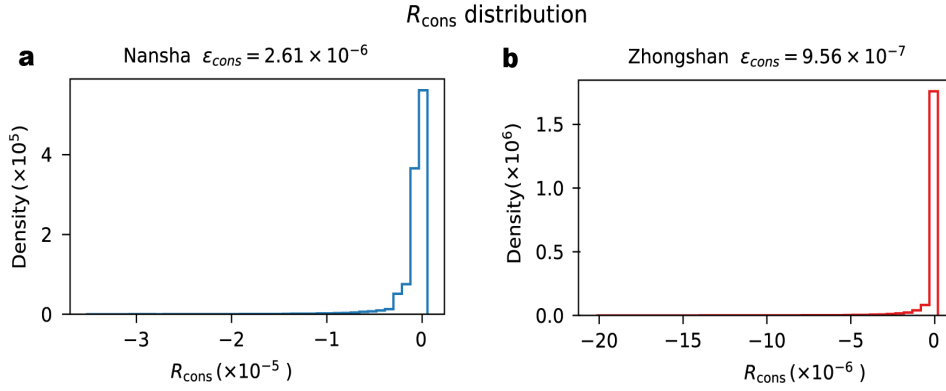


图 6-6 Distribution of the consolidation residual R_{cons} in Nansha and Zhongshan. In both cities, the residual distribution is sharply concentrated near zero, consistent with the very small values of $\varepsilon_{\text{cons}}$. This indicates that the forecast trajectories remain closely aligned with the reduced relaxation law imposed by the physics-guided framework.

laxation time τ is compared with the closure-implied time $\tau_{\text{closure}}(K, S_s, H_d)$. The mismatch is quantified by $\varepsilon_{\text{prior}}$, which remains numerically close to zero in both cities and is of order 10^{-4} in the benchmark summary. This result is important because it shows that the learned fields do not drift independently in order to absorb the data misfit. Instead, hydraulic conductivity, specific storage, drainage thickness, and emergent relaxation time form a self-consistent effective parameterisation of the observed deformation dynamics.

Together, these two diagnostics indicate that physical consistency is a verifiable property of the solution. The model does not merely produce a forecast that appears plausible. It produces trajectories that simultaneously satisfy the reduced dynamic balance and the closure structure that defines the characteristic timescale of the response. This is a substantial methodological advance over purely statistical forecasting, where physical plausibility is usually assessed only qualitatively [8, 17].

6.6.2 Spatial structure of the inferred fields

The inferred fields show that the common physical scaffold is realised differently in the two cities. In Nansha, the model recovers relatively short effective drainage paths, higher effective storage, and finer-scale spatial variability in conductivity. This pattern is consistent with a reclaimed

coastal-delta environment marked by heterogeneous fills, partially drained shallow deposits, and strong lateral contrasts [17, 10]. In Zhongshan, the inferred drainage path is generally larger and the storage field is lower over broad areas, while the conductivity and storage patterns are more spatially organised. This is consistent with a more competent and more layered hydrostratigraphic setting [40, 18].

These contrasts have a direct operational meaning. In Nansha, the physics terms help maintain interpretability and plausible uncertainty behaviour even though the data-only model can approximate the mean error closely. In Zhongshan, the physics terms contribute directly to both accuracy and uncertainty sharpness because the reduced closure aligns more tightly with the structure of the basin.

6.6.3 Identifiability and remaining trade-offs

Synthetic recovery experiments and ridge diagnostics clarify the limits of parameter interpretability. In the synthetic tests, permeability recovery through the closure is moderate, with $\text{MAE} = 0.51$ in \log_{10} units and $R^2 = 0.439$. Direct timescale recovery is weaker, with $\text{MAE} = 0.64$ in \log_{10} units and $R^2 = 0.201$. These results reinforce the interpretation advanced in Chapter 5: the most robustly recoverable quantity is the closure-consistent relaxation structure rather than the exact local decomposition into K , S_s , and H_d . These recovery patterns are shown in Figure 6-7.

The synthetic ridge analysis also shows that strong trade-off behaviour occurs only in a minority of runs, around 14%, and that the main boundary interaction is occasional contact with the upper bound on H_d . Other bounds are rarely active. This suggests that the remaining degeneracy is not primarily an artefact of clipping, but reflects a real trade-off structure among effective hydrogeological parameters in certain lithological contexts. These boundary and ridge summaries are shown in Figure 6-8. The implication is not that the inferred fields are uninformative, but that they must be interpreted as grid-scale effective fields rather than direct laboratory parameters, consistent with previous work on subsidence modelling and data assimila-

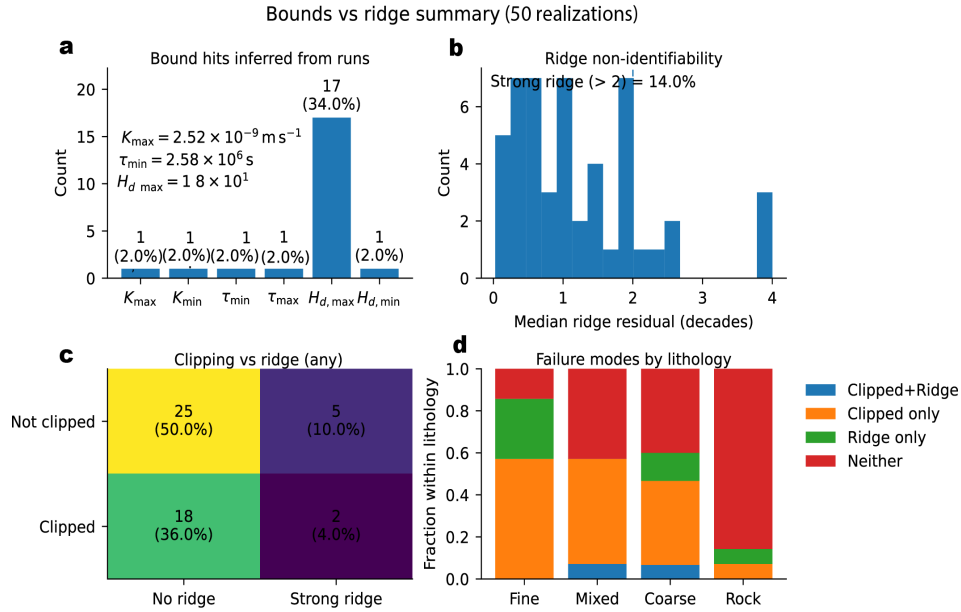


图 6-7 Synthetic identifiability experiments. Panel a shows timescale recovery, panel b permeability recovery through the closure, panel c the degeneracy-ridge diagnostic, and panel d the relation between error and the identifiability metric. The results show that the closure-consistent timescale is more robustly recoverable than the exact decomposition into conductivity, storage, and drainage thickness.

tion [20-22, 58].

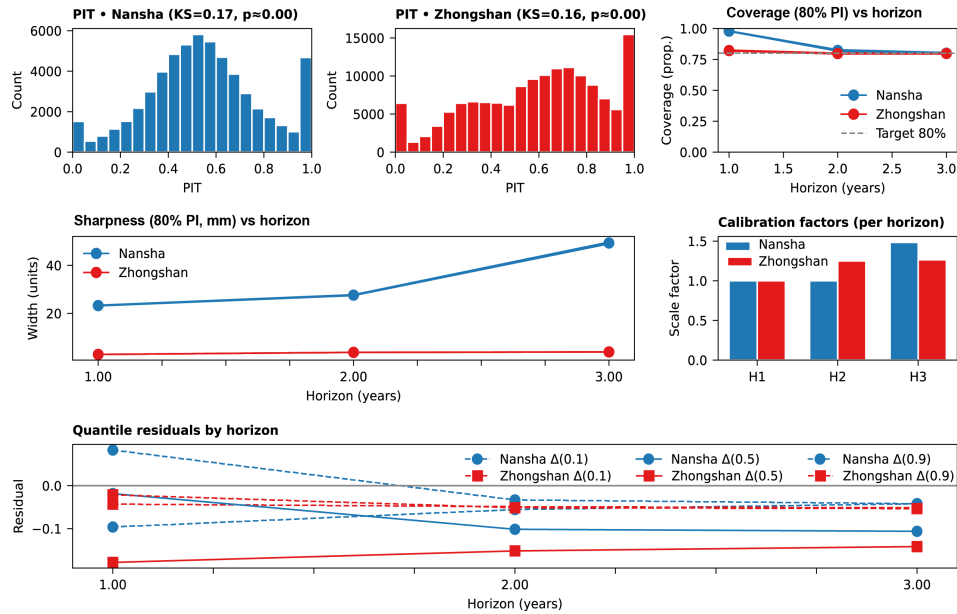


图 6-8 Bounds and ridge summary across synthetic realisations. The figure reports bound hits, the distribution of ridge non-identifiability, the relation between clipping and ridge behaviour, and the lithology-dependent distribution of failure modes. Strong ridge behaviour occurs only in a minority of realisations, whereas most solutions remain away from severe non-identifiability.

6.7 Point-support external validation

The external validation analysis provides an independent test of whether the inferred fields remain anchored to site information rather than only to the distributed deformation target. This validation was possible only in Zhongshan, where five borehole logs co-located with step-drawdown pumping tests were available. The matching distances between site data and model grid cells are small, with a median of 41.1 m and a range of 14.3–163.0 m, which supports a meaningful point-support comparison.

The strongest external result concerns the thickness pathway. The borehole-derived compressible thickness shows a positive rank association with the collocated model H_{eff} , with Spearman $\rho = 0.71$, MAE = 12.06 m, and median bias = -13.0 m. The positive rank correlation shows that the model successfully captures the relative ordering of thinner and thicker compressible sections. The negative bias indicates that the thickest sites are underestimated in absolute magnitude. This is consistent with the design of the Zhongshan thickness field, which is censor-aware and capped at 30 m. The external validation therefore supports the stratigraphic signal encoded by H_{eff} and the derived drainage pathway $H_{\text{eff}} \rightarrow H_d$, but does not imply one-to-one recovery of uncensored raw thickness in the upper tail.

The pumping tests provide a weaker check on conductivity. Using late-step specific capacity as a productivity proxy, the rank association with the collocated model K is weak (Spearman $\rho = -0.20$). This is not unexpected. Step-drawdown response depends strongly on near-well conditions, screen-specific transmissivity, and well losses, whereas the model field K is an effective cell-scale quantity embedded in a reduced closure. For that reason, the pumping-test mismatch should not be interpreted as a failure of the physics-guided framework, but as a reminder that grid-scale effective conductivity is not directly equivalent to local well productivity. These external checks are shown in Figure 6-9.

Overall, the point-support analysis independently anchors the thickness pathway that enters both the equilibrium settlement term and the clo-

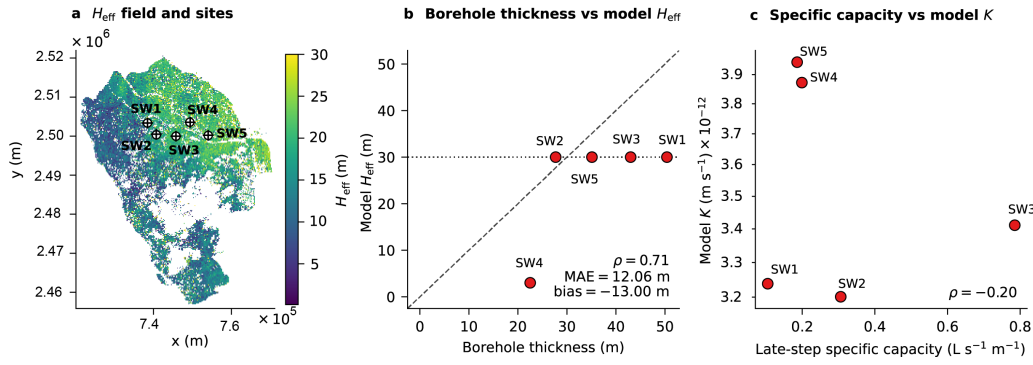


图 6-9 Point-support external validation in Zhongshan. Panel a shows the H_{eff} field and the locations of the five validation sites. Panel b compares borehole-derived compressible thickness with model H_{eff} . Panel c compares late-step specific capacity with model conductivity K . The results show a meaningful thickness-pathway agreement but a weaker relation between local pumping productivity and grid-scale effective conductivity.

sure timescale, while clarifying that sparse pumping tests are a much weaker direct constraint on effective conductivity. This is an important validation result because it identifies which parts of the inferred physics remain most robust under external checking.

6.8 Cross-city transfer and deployability

6.8.1 Retention of deterministic skill

A forecasting framework becomes substantially more useful when it can be adapted to new cities without full retraining from scratch. Transfer was therefore evaluated under three deployment modes: a target-city baseline, a direct transfer model trained on the source city and applied without adaptation, and a warm-start model initialised from the source city and then fine-tuned on the target city.

To quantify transfer, two retention statistics were used:

$$\text{Retention}_{R^2} = \frac{R^2}{R_{\text{base}}^2}, \quad \text{Retention}_{\text{MAE}} = \frac{\text{MAE}_{\text{base}}}{\text{MAE}}, \quad (6-1)$$

where the baseline is the model trained directly on the target city. Under direct zero-shot transfer, the results degrade sharply in both directions. The R^2 retention becomes negative and the MAE retention drops below 0.1, indicating order-of-magnitude increases in error relative to the target base-

line. This shows that the deformation mapping is not directly portable under strong covariate and basin structure shift.

Warm-start adaptation changes the picture substantially. After fine-tuning on limited target data, the R^2 retention returns close to unity in both transfer directions, although the MAE retention remains below one. This means that the transferable part of the model is primarily structural: the source model provides a useful initial representation of how drivers, memory, and effective physics interact, but local calibration is still needed to recover the correct amplitude of the response. This behaviour is scientifically consistent with the fact that subsidence combines shared mechanisms with basin-specific scaling linked to stratigraphy, drainage conditions, and exposure history [16, 1, 8].

6.8.2 Transfer of uncertainty and threshold risk

The uncertainty diagnostics lead to the same conclusion. Under direct transfer, coverage can collapse or become severely miscalibrated while remaining deceptively sharp. Warm-start fine-tuning moves the operating point back toward the nominal 80% interval target, sometimes with moderately wider intervals, showing that a small amount of target data is sufficient to restore uncertainty scale without discarding the source-informed prior.

Because risk management often depends on exceedance rather than raw error, the transfer study also evaluated the event $|s| \geq 50 \text{ mm yr}^{-1}$. The corresponding probability error was summarised with the Brier score,

$$\text{BS} = \frac{1}{n} \sum_{i=1}^n (p_i - o_i)^2, \quad (6-2)$$

where p_i is the predicted exceedance probability and $o_i \in \{0, 1\}$ is the observed event indicator. Direct transfer produces large Brier penalties in the more difficult direction, whereas warm-start substantially reduces them and moves the reliability curves closer to the diagonal.

6.8.3 Transfer of hotspot ranking

A final question is whether transfer preserves where the model would send an intervention team. This was evaluated by comparing the top-ranked hotspots from each transfer strategy to the target baseline using Jaccard overlap and Spearman rank correlation. Direct transfer preserves only moderate overlap. Warm-start adaptation improves both hotspot agreement and the stability of hotspot ranking. These transfer results are summarised in Figure 6-10.

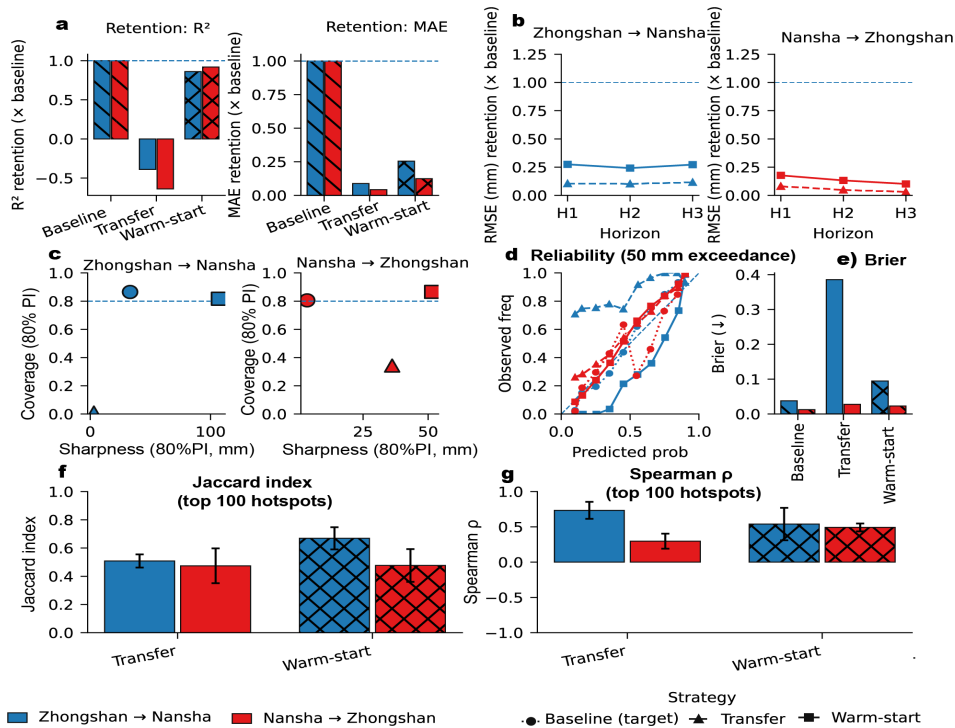


图 6-10 Cross-city transfer performance. The figure reports retention of deterministic skill, horizon-wise RMSE retention, uncertainty behaviour under transfer and warm-start adaptation, reliability for the 50 mm yr⁻¹ exceedance event, Brier score, Jaccard overlap of the top-ranked hotspots, and Spearman rank correlation of hotspot ordering. Direct zero-shot transfer is unreliable, whereas warm-start adaptation restores most of the target-city skill and substantially improves uncertainty quality.

The overall transfer result is therefore clear. GeoPriorSubsNet should not be interpreted as a universal zero-shot model for all cities. Instead, it functions as a physics-consistent initialiser whose learned representation can be adapted efficiently to new basins. This is a more realistic and more useful notion of deployability than an all-or-nothing claim of direct transfer.

6.9 Summary of the experimental evidence

The experimental results support four main conclusions.

First, the forecasting problem is genuinely regime-dependent. Nansha and Zhongshan occupy different driver spaces and different hydrogeological contexts, which explains why the same forecasting architecture operates under different uncertainty and regularisation regimes.

Second, GeoPriorSubsNet achieves strong deterministic skill in both cities and outperforms the main non-physics baselines as well as the poroelastic surrogate under the present data conditions. Its advantage is especially clear when deterministic accuracy and uncertainty quality are evaluated jointly rather than separately.

Third, the model satisfies its intended physical role. The consolidation residual remains near zero, the closure-consistency error remains numerically small, and the inferred effective fields retain physically interpretable contrasts between the two cities. The external validation supports the thickness pathway directly and clarifies the more limited interpretability of local conductivity from sparse pumping tests.

Fourth, cross-city deployment is feasible only when framed correctly. Direct transfer is unreliable under realistic distribution shift, but warm-start calibration restores usable skill and transferable risk signals. The model is therefore best viewed as a scalable, physically structured forecasting layer that can be localised efficiently rather than as a fixed universal predictor.

These results provide the empirical basis for the next chapter, which draws together the machine-learning and physics-guided stages of the study into a broader scientific discussion on mechanism, uncertainty, and the future direction of urban subsidence forecasting.

Chapter 7 Discussion

7.1 Introduction

The results presented in the preceding chapters establish a clear scientific progression. The first modelling stage showed that urban land subsidence can be simulated with strong predictive skill from heterogeneous geospatial and hydro-environmental data. The second stage then demonstrated that this predictive task can be reformulated within a physics-guided framework so that the forecast remains not only accurate, but also physically auditable and uncertainty-aware. The significance of this progression is methodological, but its implications are broader. It changes how subsidence forecasting can be understood in urban geoscience: not as a choice between flexible prediction and physical interpretation, but as a problem that requires both.

This chapter discusses the main scientific meaning of the study from six perspectives. First, it examines the transition from purely data-driven simulation to physics-guided forecasting. Second, it interprets what the results reveal about the dynamics of land subsidence in Nansha and Zhongshan. Third, it discusses the methodological contribution of the physics-guided framework. Fourth, it clarifies the meaning and limits of the inferred effective fields. Fifth, it considers the practical implications of the framework for urban risk management and basin-scale deployment. Finally, it highlights the principal limitations of the study and identifies directions for future work.

7.2 From data-driven simulation to physics-guided forecasting

The first modelling stage established that urban land subsidence can be simulated with good predictive skill from heterogeneous geospatial and hydro-environmental data. In Nansha, the machine-learning framework showed that groundwater level and building concentration carry strong explanatory power and that nonlinear models such as XGBR and LSTM can

reconstruct the observed deformation field with operationally useful accuracy [29]. This result is important because it confirms that subsidence is not beyond the reach of modern data-driven modelling, even in a rapidly urbanising coastal environment where the deformation response arises from the interaction of hydrogeology, lithology, and urban development.

At the same time, the data-driven stage made a central limitation explicit. Even when predictive skill is strong, a purely statistical model does not explain the forecast in terms of subsurface response, effective storage, drainage conditions, or compaction timescale. It can identify influential variables, but it does not provide a mechanistic account of why a predicted trajectory is physically plausible or how it should be interpreted under changing forcing. This limitation is not a minor technical issue. It concerns the scientific meaning of the forecast itself. Urban subsidence is a path-dependent deformation process rather than a flexible regression surface, and therefore requires more than correlation recovery if the objective is to support interpretation, extrapolation, and policy.

The second modelling stage addressed this limitation by introducing a physics-guided forecasting framework in which the prediction is coupled to a reduced hydro–geomechanical scaffold. This changes the role of the model. It is no longer only a device that predicts cumulative subsidence from past observations and covariates. It becomes a framework that also infers effective hydraulic conductivity, effective specific storage, and a closure-consistent relaxation timescale, while remaining compatible with a reduced consolidation balance and a groundwater residual. The forecast therefore acquires a dual meaning: it is both a probabilistic prediction and a physically auditable response of an effective urban subsurface system. This is the principal conceptual advance of the study.

This progression is justified by the experimental results. The physics-guided model preserved strong deterministic skill in both study areas, improved the quality of uncertainty in the more structured basin, and maintained near-zero consolidation and closure residuals. The methodological change therefore did not impose physical structure at the expense of predic-

tion. Rather, it showed that predictive strength and physical accountability can be combined within the same forecasting framework when the physical scaffold is kept reduced, interpretable, and compatible with the available data. This places the work in a distinct position relative to both classical numerical simulators and purely black-box machine learning [28, 30-31, 22, 58].

7.3 Scientific interpretation of the two urban subsidence regimes

The comparison between Nansha and Zhongshan provides an important scientific result in its own right. The two cities do not merely differ in error scale. They exhibit different deformation regimes. Nansha is characterised by a broader hydro-environmental range, deeper groundwater conditions, stronger heterogeneity in the inferred fields, and a more variable uncertainty structure. Zhongshan is more organised in both its driver distribution and its inferred effective fields, with smaller error, sharper prediction intervals, and a stronger direct gain from physics regularisation. These contrasts indicate that the same urban geohazard can arise through related but not identical effective dynamics across nearby basins.

The empirical driver–response behaviour supports this interpretation. In Nansha, subsidence increases strongly with depth to water over a broad range, while rainfall and urban loading mainly act as modulating effects superimposed on a highly responsive hydrogeological background. In Zhongshan, the driver range is narrower and the response structure is more compact. This suggests that Nansha is closer to a high-variability coastal compaction regime in which shallow sediments, reclamation history, and distributed anthropogenic forcing amplify local heterogeneity. Zhongshan, by contrast, behaves more like a basin where the dominant controls remain interpretable through a more stable effective structure.

The inferred effective fields are consistent with this contrast. In Nansha, the shorter inferred drainage path, stronger storage signal, and more spatially variable conductivity indicate a system with faster local adjustment but stronger spatial compartmentalisation. In Zhongshan, the larger

drainage path and more organised storage and conductivity fields indicate a slower but more structured basin response. These differences do not imply that one city is physically simple and the other complex in an absolute geological sense. Rather, they indicate that at the scale of the common annual grid, the Nansha system behaves as a more irregular effective medium, whereas Zhongshan is closer to the assumptions under which the reduced closure becomes highly informative.

This interpretation also explains why removing the physics regularisation has only a limited effect on deterministic point error in Nansha but a substantially larger effect in Zhongshan. In the more heterogeneous coastal setting, the data pathway can capture much of the empirical signal even without explicit physical structure. In the more organised setting, the reduced physical scaffold aligns more directly with the effective basin response and therefore improves both generalisation and uncertainty sharpness. This is an important result because it shows that the value of physics guidance is regime-dependent. Its contribution should be interpreted in relation to geological structure and data coherence rather than as a universal constant.

7.4 Methodological significance of the physics-guided framework

A major methodological contribution of this study lies in the way physical structure is introduced into the forecasting problem. Many previous studies place machine learning and physical modelling on opposite sides of a methodological divide: one is flexible but opaque, the other is interpretable but difficult to calibrate at scale. The present results show that this opposition is not necessary. What matters is how physical information is encoded.

The framework developed here uses a reduced closure rather than a fully resolved physical inversion. This choice is fundamental. A detailed poroelastic or basin-scale hydrogeological simulator would require boundary conditions, multi-depth loading histories, high-resolution stratigraphic continuity, and parameter information that are rarely available with suf-

efficient density across large urban domains [14, 20-21]. A purely data-driven forecast, by contrast, can be trained efficiently but leaves the internal response mechanistically opaque. The reduced closure used here occupies an intermediate position. It preserves a physically meaningful relation between conductivity, storage, drainage thickness, and characteristic timescale, while allowing the model to remain compatible with the coarse, heterogeneous, and partially censored information typical of urban geospatial archives.

The specific decision to treat H_{eff} as an externally supplied field rather than a fully free learned variable is especially important. It reduces identifiability problems, preserves the direct geological meaning of thickness, and forces the model to express the main residual flexibility through the timescale discrepancy term rather than through uncontrolled thickness drift. This design choice appears well justified by the external validation results, where borehole-derived thickness and model H_{eff} retain positive rank association despite censoring and scale mismatch. In other words, the framework is most informative when it uses the available geological information as a hard structural input rather than as another latent quantity to be absorbed into the network.

The probabilistic formulation is another important contribution. Much of the subsidence literature remains focused on deterministic prediction or susceptibility mapping [24-25, 34]. That emphasis is insufficient for urban decision-making, where the practical question is often whether a threshold will be exceeded over a planning horizon. The quantile-based structure used here allows the forecast to be evaluated in terms of coverage, sharpness, and threshold risk. More importantly, the uncertainty is not detached from the physical scaffold. The model is not trained to produce intervals in isolation; it is trained to produce intervals for trajectories that must also satisfy a reduced dynamic balance and a closure-consistent timescale. This combination is one of the distinctive methodological strengths of the work.

7.5 Interpretability, identifiability, and the meaning of the inferred fields

The inferred fields recovered by the model must be interpreted with care. The results do not support a claim of unique local inversion of hydraulic conductivity or storage from surface deformation alone. The closure relation itself makes clear why such a claim would be unrealistic. The relaxation timescale depends on a multiplicative combination of conductivity, storage, and drainage thickness, so different parameter combinations can produce very similar effective dynamics. The synthetic recovery tests confirm this point: recovery of the overall closure structure is more stable than exact recovery of each individual component.

This is not a weakness of the model in particular; it is a structural feature of the inverse problem. Subsidence observed at the land surface contains limited direct information about the unique decomposition of grid-scale hydrogeological parameters, especially in the absence of dense multi-depth head observations and detailed boundary conditions. What the model recovers most reliably is therefore the effective relaxation structure of the basin response. The learned fields K , S_s , and τ should thus be interpreted as closure-constrained effective fields, not as direct laboratory-scale or pumping-test-scale properties. This interpretation is consistent with earlier work on effective parameterisation and data assimilation in subsidence systems [20-22, 58].

Within this interpretation, the inferred fields remain scientifically valuable. They reveal spatial contrasts in effective dynamical response, they explain why two cities with the same hazard label can display different forecast behaviour, and they provide diagnostic quantities that can be audited against independent information. The point-support validation in Zhongshan is particularly informative in this respect. The positive rank agreement between borehole thickness and model H_{eff} shows that the thickness pathway remains anchored to site information, whereas the weak agreement between local step-drawdown productivity and model K reveals the expected scale mismatch between cell-scale effective conductivity and near-well test response. This is the type of interpretability that is scientifically defensible:

not a claim of literal local inversion, but a clear statement of what part of the inferred structure is robust and what part remains scale-dependent.

7.6 Practical significance for urban risk management

The practical value of the framework lies in the fact that the forecast is spatially explicit, probabilistic, and physically interpretable. Such a forecast can support decision-making in at least three ways.

The first is hotspot identification. Urban managers do not intervene everywhere at once; they prioritise locations where deformation is likely to persist or intensify. The framework developed here provides not only mean forecasts of future cumulative subsidence, but also quantile information and exceedance probabilities. This makes it possible to distinguish locations where the expected deformation is moderate but the upper-tail risk remains high from locations where both the median and the tail risk are low. Such distinctions are directly useful for ranking areas for monitoring, maintenance, or mitigation.

The second is interpretation of plausible controls. Because the model is coupled to effective hydrogeological structure, high-risk zones are not identified as purely statistical anomalies. They can be discussed in terms of stronger drawdown sensitivity, larger effective thickness, shorter or longer relaxation timescale, and stronger urban-load modulation. This type of interpretation is useful for planning because it speaks to mechanisms that can, at least in principle, be addressed by groundwater management, land-use control, or engineering adaptation.

The third is transferable deployment. The transfer experiments show that the model is not a universal zero-shot predictor. Direct source-to-target application fails under significant covariate and basin-structure shift. However, warm-start fine-tuning recovers most of the baseline skill and substantially improves threshold-risk reliability. This is a practically important outcome. It means that the framework can be used as a physics-consistent initial layer for new cities, with modest local data being sufficient to recover operational skill. In urban forecasting, this form of transfer is often more

useful than the unrealistic goal of perfect zero-shot generalisation.

These practical implications align with the broader literature on subsidence, flood exposure, and urban resilience [9, 11, 10]. Land subsidence is not only a geological process; it is an amplifier of urban risk. A forecasting framework that helps locate persistent hotspots, quantify uncertainty, and preserve interpretable physical structure can therefore contribute meaningfully to sustainable groundwater management and long-term infrastructure planning.

7.7 Limitations of the study

Several limitations should be stated clearly.

The first concerns scale. The modelling framework operates on a common annual grid and therefore represents basin response at an effective urban scale rather than at the scale of individual wells, individual structures, or detailed stratigraphic layers. This is appropriate for city-scale forecasting, but it necessarily smooths processes that may be important for local engineering analysis. The inferred fields should therefore not be interpreted as substitutes for site-specific hydrogeological characterisation.

The second concerns model structure. The reduced consolidation law and closure-consistent timescale provide a useful scaffold, but they remain simplifications. Real urban subsidence may involve multi-layer delayed drainage, anisotropy, lateral boundary effects, non-linear constitutive behaviour, and engineered perturbations that are not fully represented in the present closure. The good performance of the model does not mean that these effects are absent; it means that the reduced structure is sufficient to organise a large fraction of the observed basin-scale response.

The third concerns data limitations. The thickness field in Zhongshan is censor-aware and capped in the upper tail, which constrains the degree to which absolute thickness magnitude can be recovered or externally validated. Groundwater information is also limited by the spatial and temporal density of available observations. Rainfall is included as a future-known

forcing, but future pumping and future urban loading are not explicitly prescribed as scenario variables in the present framework. This limits the analysis to forecasting under the continuation or statistical extension of the observed driver system rather than under fully user-defined policy scenarios.

The fourth concerns identifiability. As discussed above, the closure permits non-unique trade-offs between conductivity, storage, and drainage thickness. The framework mitigates this problem through priors, bounds, and an externally supplied thickness field, but it does not eliminate the structural ambiguity of the inverse problem. This is why the inferred quantities are interpreted as effective fields and why external validation is focused more strongly on thickness than on local conductivity.

The fifth concerns transfer. Warm-start adaptation performs well, but direct transfer remains poor. This implies that basin-specific calibration is still necessary and that the framework should be viewed as a transferable structured prior rather than a fully universal model.

7.8 Future research directions

Several directions follow naturally from these results.

A first priority is to extend the framework toward richer physical representation without sacrificing scalability. One possible path is to move from a single effective drainage thickness toward a multi-layer or depth-aware structure in which separate delay pathways can be resolved for shallow and deep compressible units. Such an extension would make the physical interpretation of timescale more specific while preserving the advantages of reduced-order learning.

A second priority is to incorporate richer scenario forcing. The present framework handles rainfall as a future-known variable, but future groundwater extraction, planned construction intensity, or engineered dewatering could be introduced as explicit scenario controls. This would move the model closer to a decision-support instrument in which users can explore management choices rather than only extrapolate observed dynamics.

A third direction concerns data assimilation and multimodal constraints. The current study already shows the value of combining InSAR, groundwater, borehole thickness, and urban-development information. Additional constraints from GNSS, levelling, pumping-test archives, or even foundation inventories could help reduce uncertainty in the inferred effective fields, especially in basins where thickness and conductivity remain weakly identifiable from deformation alone.

A fourth direction concerns transfer learning across larger basin families. The present comparison uses two study areas within the Pearl River Delta. Extending the framework across more diverse geological and climatic settings would help identify which aspects of the learned representation are broadly transferable and which remain strongly regional. This would be an important step toward a genuinely scalable urban subsidence forecasting framework.

A fifth direction concerns calibration diagnostics. The results already show that nominal coverage and interval sharpness can be evaluated in a physically meaningful way. Further work could refine this by combining probabilistic scoring rules, horizon-specific risk metrics, and more formal calibration correction under distribution shift.

7.9 Conclusion and perspectives

The principal contribution of this study is not simply the construction of a more accurate model. Its broader significance lies in the fact that it reframes urban subsidence forecasting as a problem that requires predictive strength, uncertainty awareness, and physical interpretability simultaneously. The machine-learning stage established that heterogeneous urban data contain sufficient signal to support useful forecasting. The physics-guided stage then showed that this forecasting problem can be reformulated so that the trajectories remain auditable against a reduced hydro-geomechanical scaffold and that the resulting model retains strong performance across two contrasting subsidence regimes.

This progression suggests a broader lesson for geoscientific modelling.

In problems such as urban land subsidence, the most useful models may not be those that maximise physical detail at all cost or those that maximise predictive flexibility in isolation. The most useful models may instead be those that preserve the essential physical structure of the system while remaining compatible with the scale, incompleteness, and heterogeneity of the available data. That is the position reached by the present study, and it defines the foundation on which future work can build.

参考文献

- [1] CHAUSSARD E, WADOWINSKI S, CABRAL-CANO E, et al. Land subsidence in central mexico detected by ALOS InSAR time-series [J/OL]. *Remote Sensing of Environment*, 2014, 140: 94-106. DOI: [10.1016/j.rse.2013.08.038](https://doi.org/10.1016/j.rse.2013.08.038).
- [2] PACHECO-MARTINEZ J, HERNANDEZ-MARIN M, BURBEY T J, et al. Land subsidence and ground failure associated to groundwater exploitation in the aguascalientes valley, México[J]. *Engineering Geology*, 2013, 164: 172-186.
- [3] PHIEN-WEJ N, GIAO P H, NUTALAYA P. Land subsidence in bangkok, thailand[J/OL]. *Engineering Geology*, 2006, 82(4): 187-201. DOI: [10.1016/j.enggeo.2005.10.004](https://doi.org/10.1016/j.enggeo.2005.10.004).
- [4] VAN NIEKERK H J, DER WALT I J. Dewatering of the far west rand dolomitic area by gold mining activities and subsequent ground instability[J]. *Land Degradation & Development*, 2006, 17: 441-452.
- [5] YIN J, YU D, YIN Z, et al. Evaluating the impact and risk of pluvial flash flood on intra-urban road network: a case study in the city center of shanghai, china[J]. *Journal of Hydrology*, 2016, 537: 138-145.
- [6] ANDERSON J F. Geological subsidence and upheaval[J/OL]. *Nature*, 1873, 8(194): 223-224. <https://doi.org/10.1038/008223c0>.
- [7] BAGHERI-GAVKOSH M, HOSSEINI S M, ATAIE-ASHTIANI B, et al. Land subsidence: a global challenge[J/OL]. *Science of the Total Environment*, 2021, 778. DOI: [10.1016/j.scitotenv.2021.146193](https://doi.org/10.1016/j.scitotenv.2021.146193).
- [8] SHIRZAEI M, FREYMUELLER J, TÖRNQVIST T E, et al. Measuring, modelling and projecting coastal land subsidence[J/OL]. *Nature Reviews Earth & Environment*, 2021, 2(1): 40-58. <https://doi.org/10.1038/s43017-020-00115-x>.

- [9] NICHOLLS R J, LINCKE D, HINKEL J, et al. A global analysis of subsidence, relative sea-level change and coastal flood exposure[J]. *Nature Climate Change*, 2021, 11(4): 338-342.
- [10] HASAN M F, SMITH R, VAJEDIAN S, et al. Global land subsidence mapping reveals widespread loss of aquifer storage capacity[J/OL]. *Nature Communications*, 2023, 14(1): 1-10. DOI: [10.1038/s41467-023-41933-z](https://doi.org/10.1038/s41467-023-41933-z).
- [11] FANG J, NICHOLLS R J, BROWN S, et al. Benefits of subsidence control for coastal flooding in china[J/OL]. *Nature Communications*, 2022, 13(1). DOI: [10.1038/s41467-022-34525-w](https://doi.org/10.1038/s41467-022-34525-w).
- [12] GALLOWAY D L, BURBEY T J. Regional land subsidence accompanying groundwater extraction[J/OL]. *Hydrogeology Journal*, 2011, 19(8): 1459. DOI: [10.1007/s10040-011-0775-5](https://doi.org/10.1007/s10040-011-0775-5).
- [13] BUDHU M, ADIYAMAN I B. Mechanics of land subsidence due to groundwater pumping[J]. *International Journal for Numerical and Analytical Methods in Geomechanics*, 2010, 34: 1459-1478.
- [14] HOFFMANN J, LEAKE S A, GALLOWAY D L, et al. MODFLOW-2000 ground-water model — user guide to the subsidence and aquifer-system compaction (SUB) package[M/OL]. 2000. <https://pubs.usgs.gov/of/2003/ofr03-233/pdf/ofr03233.pdf>.
- [15] XUE Y Q, ZHANG Y, YE S J, et al. Land subsidence in china[J/OL]. *Environmental Geology*, 2005, 48: 713-720. DOI: [10.1007/s00254-005-0010-6](https://doi.org/10.1007/s00254-005-0010-6).
- [16] ABIDIN H Z, ANDREAS H, GUMILAR I, et al. Land subsidence of jakarta (Indonesia) and its relation with urban development[J/OL]. *Natural Hazards*, 2011, 59(3): 1753-1771. DOI: [10.1007/s11069-011-9866-9](https://doi.org/10.1007/s11069-011-9866-9).
- [17] CANDELA T, KOSTER K. The many faces of anthropogenic subsidence[J/OL]. *Science*, 2022, 376(6600): 1381-1382. DOI: [10.1126/science.abn3676](https://doi.org/10.1126/science.abn3676).

- [18] CHEN B, GONG H, CHEN Y, et al. Land subsidence and its relation with groundwater aquifers in beijing plain of china[J]. *Science of the Total Environment*, 2020, 735: 139111.
- [19] CIGNA F, TAPETE D. Urban growth and land subsidence: multi-decadal investigation using human settlement data and satellite InSAR in morelia, mexico[J/OL]. *Science of the Total Environment*, 2022, 811: 152211. <https://doi.org/10.1016/j.scitotenv.2021.152211>.
- [20] BA D, ALZRAIEE A, ZOCCARATO C, et al. Testing a data assimilation approach to reduce geomechanical uncertainties in modelling land subsidence[J/OL]. *Environmental Geotechnics*, 2014, 3(6): 386-396. <https://www.sciencedirect.com/science/article/pii/S2051803X14000171>. DOI: 10.1680/envgeo.15.00005.
- [21] GAZZOLA L, FERRONATO M, FRIGO M, et al. A novel methodological approach for land subsidence prediction through data assimilation techniques[J]. *Computational Geosciences*, 2021, 25: 1731-1750.
- [22] VERBERNE M, KOSTER K, LOURENS A, et al. Disentangling shallow subsidence sources by data assimilation in a reclaimed urbanized coastal plain, south flevoland polder, the netherlands[J/OL]. *Journal of Geophysical Research: Earth Surface*, 2023, 128: 1-23. DOI: [10.1029/2022JF007031](https://doi.org/10.1029/2022JF007031).
- [23] ZAPATA-NORBERTO B, MORALES-CASIQUE E, HERRERA G S. One-dimensional simulation of land subsidence in vertically-heterogeneous highly compressible aquitards coupled with data assimilation via ensemble kalman filter[J/OL]. *Environmental Modelling & Software*, 2025, 194: 106690. DOI: [10.1016/j.envsoft.2025.106690](https://doi.org/10.1016/j.envsoft.2025.106690).
- [24] RAHMATI O, GOLKARIAN A, BIGGS T, et al. Land subsidence hazard modeling: machine learning to identify predictors and the role of human activities[J/OL]. *Journal of Environmental Management*, 2019, 236(August 2018): 466-480. <https://doi.org/10.1016/j.jenvman.2019.02.020>.

- [25] WANG Z, GUO L, GONG H, et al. Land subsidence simulation based on extremely randomized trees combined with monte carlo algorithm [J/OL]. *Computers and Geosciences*, 2023, 178(July): 105415. <https://doi.org/10.1016/j.cageo.2023.105415>.
- [26] LI H, ZHU L, DAI Z, et al. Spatiotemporal modeling of land subsidence using a geographically weighted deep learning method based on PS-InSAR[J]. *Science of the Total Environment*, 2021, 799: 149244.
- [27] SHI L, GONG H, CHEN B, et al. Land subsidence prediction induced by multiple factors using machine learning method[J/OL]. *Remote Sensing*, 2020, 12: 4044. DOI: [10.3390/rs12244044](https://doi.org/10.3390/rs12244044).
- [28] GUZY A, MALINOWSKA A A. State of the art and recent advancements in the modelling of land subsidence induced by groundwater withdrawal[J]. *Water*, 2020, 12: 2051.
- [29] LIU J, LIU W, ALLECHY F B, et al. Machine learning-based techniques for land subsidence simulation in an urban area[J/OL]. *Journal of Environmental Management*, 2024, 352: 120078. DOI: [10.1016/j.jenvman.2024.120078](https://doi.org/10.1016/j.jenvman.2024.120078).
- [30] DONNELLY J, DANESHKHAH A, ABOLFATHI S. Physics-informed neural networks as surrogate models of hydrodynamic simulators[J/OL]. *Science of The Total Environment*, 2023, 912: 168814. DOI: [10.1016/j.scitotenv.2023.168814](https://doi.org/10.1016/j.scitotenv.2023.168814).
- [31] ROY S, ANNAVARAPU C, ROY P, et al. Physics-informed neural networks for heterogeneous poroelastic media[J/OL]. *International Journal for Computational Methods in Engineering Science and Mechanics*, 2024, 26(2): 187-207. DOI: [10.1080/15502287.2024.2440420](https://doi.org/10.1080/15502287.2024.2440420).
- [32] SARMA A K, ROY S, ANNAVARAPU C, et al. Interface pinns (i-pinns): a physics-informed neural networks framework for interface problems[J/OL]. *Computer Methods in Applied Mechanics and Engineering*, 2024, 429: 117135. DOI: [10.1016/j.cma.2024.117135](https://doi.org/10.1016/j.cma.2024.117135).

- [33] ROY S, ANNAVARAPU C, ROY P, et al. Physics-informed neural networks for heterogeneous poroelastic media[A/OL]. arXiv (2024). <https://arxiv.org/abs/2407.03372>.
- [34] ZAMANIRAD M, SARRAF A, SEDGHI H, et al. Modeling the influence of groundwater exploitation on land subsidence susceptibility using machine learning algorithms[J/OL]. Natural Resources Research, 2020, 29: 1127-1141. DOI: [10.1007/s11053-019-09490-9](https://doi.org/10.1007/s11053-019-09490-9).
- [35] LIM B, ARIK S O, LOEFF N, et al. Temporal fusion transformers for interpretable multi-horizon time series forecasting[J/OL]. International Journal of Forecasting, 2021, 37(4): 1748-1764. DOI: [10.1016/j.ijforecast.2021.03.012](https://doi.org/10.1016/j.ijforecast.2021.03.012).
- [36] KOUADIO K L, LIU Z, LIU R, et al. XTFT: a next-generation temporal fusion transformer for uncertainty-rich time series forecasting [EB/OL]. 2025. <https://authorea.com/users/643438/articles/711823-xtft-a-next-generation-temporal-fusion-transformer-for-uncertainty-rich-time-series-forecasting>. DOI: [10.22541/au.175390529.91420978.v1](https://doi.org/10.22541/au.175390529.91420978.v1).
- [37] United Nations. Sustainable development: the 17 goals[EB/OL]. 2020. <https://sdgs.un.org/goals>.
- [38] United Nations. Global sustainable development report 2023: times of crisis, times of change: science for accelerating transformations to sustainable development[EB/OL]. 2023. https://sdgs.un.org/sites/default/files/2023-09/FINAL%20GSDR%202023-Digital-110923_1.pdf.
- [39] FERNÁNDEZ-TORRES E A, CABRAL-CANO E, NOVELO-CASANOVA D A, et al. Risk assessment of land subsidence and associated faulting in Mexico City using InSAR[J]. Natural Hazards, 2022, 112(1): 37-55.
- [40] ZHU L, GONG H, LI X, et al. Land subsidence due to groundwater withdrawal in the northern Beijing plain, China[J/OL]. Engineering

- Geology, 2015, 193: 243-255. <http://dx.doi.org/10.1016/j.enggeo.2015.04.020>.
- [41] LI Z, CHEN Q, XUE Y, et al. Numerical investigation of processes, features, and control of land subsidence caused by groundwater extraction and coal mining: a case study from eastern china[J]. Environmental Earth Sciences, 2023, 82: 82.
- [42] STROZZI T, WEGMULLER U, TOSI L, et al. Land subsidence monitoring with differential SAR interferometry[J]. Photogrammetric Engineering and Remote Sensing, 2001, 67(11): 1261-1270.
- [43] SCHMIDT D A, BÜRGMANN R. Time-dependent land uplift and subsidence in the santa clara valley, california, from a large interferometric synthetic aperture radar data set[J/OL]. Journal of Geophysical Research: Solid Earth, 2003, 108(B9). DOI: [10.1029/2002jb002267](https://doi.org/10.1029/2002jb002267).
- [44] TIZZANI P, BERARDINO P, CASU F, et al. Surface deformation of long valley caldera and mono basin, california, investigated with the SBAS-InSAR approach[J]. Remote Sensing of Environment, 2007, 108: 277-289.
- [45] FOROUGHNIA F, NEMATI S, MAGHSOUDI Y, et al. An iterative PS-InSAR method for the analysis of large spatio-temporal baseline data stacks for land subsidence estimation[J]. International Journal of Applied Earth Observation and Geoinformation, 2019, 74: 248-258.
- [46] BAI L, JIANG L, WANG H, et al. Spatiotemporal characterization of land subsidence and uplift (2009–2010) over wuhan in central china revealed by TerraSAR-X InSAR analysis[J]. Remote Sensing, 2016, 8: 350.
- [47] SUN M, DU Y, LIU Q, et al. Understanding the spatial-temporal characteristics of land subsidence in shenzhen under rapid urbanization based on MT-InSAR[J/OL]. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 2023, PP: 1-13. DOI: [10.1109/JSTARS.2023.3264652](https://doi.org/10.1109/JSTARS.2023.3264652).

- [48] ZHOU C, LAN H, GONG H, et al. Reduced rate of land subsidence since 2016 in beijing, china: evidence from tomo-PSInSAR using RadarSAT-2 and sentinel-1 datasets[J]. *International Journal of Remote Sensing*, 2020, 41: 1259-1285.
- [49] MARFAI M A, KING L. Monitoring land subsidence in semarang, indonesia[J/OL]. *Environmental Geology*, 2007, 53: 651-659. DOI: [10.1007/s00254-007-0680-3](https://doi.org/10.1007/s00254-007-0680-3).
- [50] ZHAN Y, ZHANG Y, ZHANG J, et al. Risk assessment of land subsidence in shanghai municipality based on AHP and EWM[J/OL]. *Scientific Reports*, 2025, 15(1): 1-19. DOI: [10.1038/s41598-025-91109-6](https://doi.org/10.1038/s41598-025-91109-6).
- [51] DU Z, GE L, NG A H M, et al. Analysis of the impact of the south-to-north water diversion project on water balance and land subsidence in beijing, china between 2007 and 2020[J]. *Journal of Hydrology*, 2021, 603: 126990.
- [52] PAN Y, DING H, LI J, et al. Transient hydrology-induced elastic deformation and land subsidence in australia constrained by contemporary geodetic measurements[J]. *Earth and Planetary Science Letters*, 2022, 588: 117556.
- [53] NISHI K, KAWAI M, CAHYONO B E, et al. Consecutive DInSAR and well based on the law of material conservation between land surface pressure and ground water to observe land subsidence[J]. *Geocarto International*, 2023, 38: 2159069.
- [54] SHAHBAZI S, MOUSAVI Z, REZAEI A. Constraints on the hydrogeological properties and land subsidence through GNSS and InSAR measurements and well data in salmas plain, northwest of urmia lake, iran[J]. *Hydrogeology Journal*, 2022, 30: 533-555.
- [55] RAJU A, NANDA R, SINGH A, et al. Multi-temporal analysis of groundwater depletion-induced land subsidence in central ganga al-

- luvial plain, northern india[J]. *Geocarto International*, 2022, 37(26): 11732-11755.
- [56] BENNETHUM L S, MURAD M A, CUSHMAN J H. Modified darcy's law, terzaghi's effective stress principle and fick's law for swelling clay soils[J/OL]. *Computers and Geotechnics*, 1997, 20(3): 245-266. DOI: [10.1016/S0266-352X\(97\)00005-0](https://doi.org/10.1016/S0266-352X(97)00005-0).
- [57] MAHMOUDPOUR M, KHAMEHCHIYAN M, NIKUDEL M R, et al. Numerical simulation and prediction of regional land subsidence caused by groundwater exploitation in the southwest plain of tehran, iran[J/OL]. *Engineering Geology*, 2016, 201(2016): 6-28. <http://dx.doi.org/10.1016/j.enggeo.2015.12.004>.
- [58] VERBERNE M, TEATINI P, KOSTER K, et al. An integral approach using InSAR and data assimilation to disentangle and quantify multi-depth driven subsidence causes in the ravenna coastland, northern italy [J/OL]. *Geomechanics for Energy and the Environment*, 2025, 43 (May): 100710. <https://doi.org/10.1016/j.gete.2025.100710>.
- [59] EBRAHIMY H, FEIZIZADEH B, SALMANI S, et al. A comparative study of land subsidence susceptibility mapping of tasuj plane, iran, using boosted regression tree, random forest and classification and regression tree methods[J]. *Environmental Earth Sciences*, 2020, 79: 1-12.
- [60] MOHAMMADY M, POURGHASEMI H R, AMIRI M. Assessment of land subsidence susceptibility in semnan plain (Iran): a comparison of support vector machine and weights of evidence data mining algorithms[J]. *Natural Hazards*, 2019, 99: 951-971.
- [61] DENG Z, KE Y, GONG H, et al. Land subsidence prediction in beijing based on PS-InSAR technique and improved grey-markov model[J]. *GIScience & Remote Sensing*, 2017, 54: 797-818.
- [62] ZHOU C, GONG H, CHEN B, et al. Quantifying the contribution

- of multiple factors to land subsidence in the beijing plain, china with machine learning technology[J]. *Geomorphology*, 2019, 335: 48-61.
- [63] FRIEDMAN J H. Greedy function approximation: a gradient boosting machine[J/OL]. *The Annals of Statistics*, 2001, 29: 1189-1232. DOI: [10.2307/2699986](https://doi.org/10.2307/2699986).
- [64] GERS F A, SCHMIDHUBER J, CUMMINS F. Learning to forget: continual prediction with LSTM[J]. *Neural Computation*, 2000, 12: 2451-2471.
- [65] KRATZERT F, KLOTZ D, BRENNER C, et al. Rainfall–runoff modelling using long short-term memory (LSTM) networks[J]. *Hydrology and Earth System Sciences*, 2018, 22: 6005-6022.
- [66] KOCH J, SCHNEIDER R. Long short-term memory networks enhance rainfall-runoff modelling at the national scale of denmark[J]. *GEUS Bulletin*, 2022, 49.
- [67] RAHMATI O, FALAH F, NAGHIBI S A, et al. Land subsidence modelling using tree-based machine learning algorithms[J/OL]. *Science of the Total Environment*, 2019, 672: 239-252. DOI: [10.1016/j.scitotenv.2019.03.496](https://doi.org/10.1016/j.scitotenv.2019.03.496).
- [68] HO T K. Random decision forests[C]//*Proceedings of 3rd International Conference on Document Analysis and Recognition*. 1995: 278-282.
- [69] NOURANI V, ELKIRAN G, ABDULLAHI J. Multi-station artificial intelligence based ensemble modeling of reference evapotranspiration using pan evaporation measurements[J]. *Journal of Hydrology*, 2019, 577: 123958.
- [70] RAJAEI T, EBRAHIMI H, NOURANI V. A review of the artificial intelligence methods in groundwater level modeling[J]. *Journal of Hydrology*, 2019, 572: 336-351.

- [71] TAYLOR K E. Summarizing multiple aspects of model performance in a single diagram[J/OL]. *Journal of Geophysical Research: Atmospheres*, 2001, 106(D7): 7183-7192. DOI: [10.1029/2000JD900719](https://doi.org/10.1029/2000JD900719).

Achievements

Published papers

1. **Kouadio, Kouao Laurent.** k-diagram: Rethinking forecasting uncertainty via polar-based visualization [J]. *Journal of Open Source Software*, 2025, 10(116): 8661. DOI: <https://doi.org/10.21105/joss.08661>.
2. **Kouadio K. L.**, Liu J., Liu W., and Liu R. A mixture learning strategy for predicting aquifer permeability coefficient K [J]. *Computers & Geosciences*, 2025. DOI: <https://doi.org/10.1016/j.cageo.2024.105819>.
3. **Kouadio K. L.** and Bizimana P. C. Adaptive Hammerstein–Wiener Model for Intelligent Networks [C]. *2024 International Conference on Intelligent Computing and Next Generation Networks (ICNGN)*, 2024. DOI: <https://doi.org/10.1109/ICNGN63705.2024.10871420>.
4. Liu R., Jiang S., Ou J., **Kouadio K. L.**[†], and Xiong B. Multifaceted anomaly detection framework for leachate monitoring in landfills [J]. *Journal of Environmental Management*, 2024. DOI: <https://doi.org/10.1016/j.jenvman.2024.122130>.
5. **Kouadio K. L.**, Liu J., Liu W., Liu R., and Boukhalifa Z. An integrated approach for sewage diversion: Case of Huayuan mine, Hunan Province, China [J]. *Geophysics*, 2024. DOI: <https://doi.org/10.1190/geo2023-0332.1>.
6. **Kouadio K. L.**, Liu J., Liu R., Wang Y., and Liu W. K-Means Featurizer: A booster for intricate datasets [J]. *Earth Science Informatics*, 2024. DOI: <https://doi.org/10.1007/s12145-024-01236-3>.
7. Liu J., Liu W., Allechy F. B., Zheng Z., Liu R., and **Kouadio K. L.**[†]. Machine learning-based techniques for land subsidence simulation in an urban area [J]. *Journal of Environmental Management*, 2024, 352: 120078. DOI: <https://doi.org/10.1016/j.jenvman.2024.120078>.

8. **Kouadio K. L.**, Liu J., Kouamelan S. K., and Liu R. Ensemble learning paradigms for flow rate prediction boosting [J]. *Water Resources Management*, 2023, 37: 4413–4431. DOI: <https://doi.org/10.1007/s11269-023-03562-5>.
9. **Kouadio K. L.**, Liu J., and Liu R. watex: machine learning research in water exploration [J]. *SoftwareX*, 2023, 22: 101367. DOI: <https://doi.org/10.1016/j.softx.2023.101367>.
10. **Kouadio K. L.**, Liu R., Malory A. O., Liu W., and Liu C. A novel approach for water reservoir mapping using controlled-source audio-frequency magnetotelluric in Xingning area, Hunan Province, China [J]. *Geophysical Prospecting*, 2023. DOI: <https://doi.org/10.1111/1365-2478.13385>.

† Corresponding author.

Manuscripts under review

11. **Kouadio K. L.**, Liu R., Jiang S., Liu Z., Kouamelan S., Liu W., Qing Z., and Zheng Z. Physics-Informed Deep Learning Reveals Divergent Urban Land Subsidence Regimes. Submitted to *Nature Geosciences*, 2026.
12. **Kouadio K. L.**, Liu R., Loukou K. G. H., Liu W., Qing Z., and Liu Z. A Diagnostic Framework for Spatiotemporal Forecast Uncertainty. Submitted to *Environmental Modelling & Software*, 2025.
13. **Kouadio K. L.**, Liu Z., Liu R., Bizimana P. C., Yang G., and Liu W. XTFT: A Next-Generation Temporal Fusion Transformer for Uncertainty-Rich Time Series Forecasting. Under review at *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. DOI: <https://doi.org/10.22541/au.175390529.91420978.v1>.

Acknowledgements

First and foremost, I would like to express my sincere and profound gratitude to Professor Jianxin Liu, my postdoctoral supervisor and mentor, for giving me the opportunity to join Central South University as a postdoctoral fellow. His trust, guidance, scientific vision, and continuous support have been fundamental to the development of this work. I am deeply grateful for his encouragement, his high academic standards, and the stimulating research environment he has provided through the Hunan Key Laboratory of Nonferrous Resources and Geological Hazards Exploration. His support has been decisive in shaping both the direction and quality of this research journey.

I would also like to express my heartfelt thanks to Dr. Rong Liu for her valuable guidance, collaboration, and constant encouragement throughout my postdoctoral research. Her insight, rigor, and support greatly strengthened the scientific development of this work.

My sincere appreciation also goes to Professor Jianghai Xia, my Ph.D. advisor, whose mentorship laid the intellectual foundation of my work in geophysics, computational science, and the application of machine learning to environmental and geophysical problems. I remain deeply thankful for his guidance and for the lasting influence he has had on my scientific formation.

I am equally grateful to Dr. Kouamelan Serge Kouamelan, Dr. Francis Oppong, and Dr. Moutanou Modeste Judes Zeye for their support, encouragement, collaboration, and valuable discussions. Their different perspectives, spanning geophysics, computational modelling, engineering, and scientific software, have enriched my research and broadened its scope.

I would like to acknowledge with gratitude the support of the National Natural Science Foundation of China, whose funding under grant numbers 42130810, 42004065, 42074165, 41774149, and 42374093 made important parts of this research possible, including the studies related to Huayuan mine and urban land subsidence.

I am particularly thankful to the Guangdong Provincial Geological Bureau for providing hydrogeology and engineering geology data, as well as ground subsidence measurement data that were essential for the land-subsidence studies presented in this report.

I also gratefully acknowledge Ibrahim Berthe and Kone Saramatou of the Water Resources Department, and Diarassouba Mada of the Côte d'Ivoire National Office for Drinking Water (ONEP), for their collaboration and support during the PPU and PNAEP programs. My sincere thanks are also extended to Sombo Boko Celestin and Gahe Emile of the School of Earth Science and Mineral Resources (STRM) for helping to establish the monitoring and exchange framework during the second phase of the project. I am grateful to Dankoh Gnamin for providing the geological map of the experimental area, and to Ndeutala Selma Iita for helping improve the readability of the report.

I would further like to thank the editors and anonymous reviewers of my published and submitted papers. Their careful reading, constructive criticism, and insightful comments have significantly improved the quality, clarity, and scientific rigor of this body of work.

My appreciation also goes to all colleagues, collaborators, faculty members, and administrative staff at Central South University who contributed, directly or indirectly, to making this postdoctoral period both productive and enriching.

Finally, I would like to express my deepest gratitude to my family for their unconditional love, patience, support, and encouragement throughout this journey. Their faith in me has remained a constant source of strength.

To all those who have accompanied me, supported me, and believed in this work, I offer my sincere thanks.

KOUADIO Kouao Laurent

2026