International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 1
Article Received: *25 December 2025* ***Revised:*** *12 January 2026* ***Accepted:*** *02 February 2026*

AI-Driven Regulation-Aware BIM Framework for Clash Prioritisation and Digital Twin Integration in Saudi Vision 2030 Megaprojects

Hussam Hesham Zakieh

Master Degree- Master in Engineering Project Management, Saudi Arabia / Al Riyadh / 05010

Email: hussamzakiah@gmail.com , ORCID: https://orcid.org/0009-0000-4934-7361

Email: h.zakieh@sak-consult.com

ABSTRACT

Purpose
This study develops an AI-driven, regulation-aware framework that integrates Building Information Modelling (BIM) and the Saudi Building Code (SBC) to prioritise critical clashes and enable digital twin integration within Vision 2030 megaprojects.

Design/methodology/approach
A hybrid ensemble combining Random Forest (RF), Convolutional Neural Network (CNN), Graph Convolutional Network (GCN), and Graph Attention Network (GAT) was trained using hierarchical graph processing. SBC clauses were encoded into IFC features, with calibrated probabilities and a fixed cost-derived decision threshold.

Findings
Nested leave-one-project-out (LOPO) testing across five industrial federations demonstrated consistent improvements in AUROC, AUPRC, and calibration. The framework reduced coordination time by approximately 65% compared with incumbent workflows, with statistically significant results and large effect sizes.

Originality/value
This paper presents one of the first regulation-aware AI models for BIM clash prioritisation under the Saudi Building Code. The openly released framework enables reproducibility and provides a foundation for real-time digital twins in Saudi Vision 2030 projects.

Keywords: Digital Twin; Vision 2030; BIM; Clash Prioritisation; Saudi Building Code; IFC; Graph Neural Networks; Calibration; Cost-sensitive classification; Automation in Construction

1. INTRODUCTION

Giga-scale developments in Saudi Arabia (e.g., the PIF portfolio) compel highly coordinated, code-aligned design at unprecedented rates. Commercial clash engines excel at geometric *detection* but still inundate coordinators with thousands of low-consequence warnings, obscuring code-sensitive conflicts that should be prioritised for triage. Within Saudi Arabia's Vision 2030 transformation agenda, giga-scale developments such as NEOM, Qiddiya, and the Red Sea Project exemplify an unprecedented concentration of digitally coordinated construction. These programs prioritise sustainability, safety, and automation through the convergence of Artificial Intelligence (AI), BIM, and digital twin technologies. However, the ability to translate regulatory provisions—particularly the Saudi Building Code (SBC)—into actionable intelligence for automated quality control remains limited. Addressing this gap requires AI frameworks that integrate code semantics with data-driven reasoning directly within BIM environments. This paper advocates for a shift from mere detection to regulation-aware prioritisation, wherein criticality is modelled based on geometry, semantics, *and* explicit obligations encoded from the Saudi Building Code (SBC) into IFC-aligned features [1–5].

Accordingly, this study contributes to Vision 2030's digital construction agenda through three advances. First,

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 1
Article Received: *25 December 2025* ***Revised:*** *12 January 2026* ***Accepted:*** *02 February 2026*

we formalise a pipeline that (i) parses IFC, (ii) maps BSBC clauses to verifiable features, and (iii) learns a hybrid ensemble RF (tabular), CN (clash-centric views), GCN/GAT (relational context) with hierarchical graph processing (HGP) for scalability. Second, we calibrate probabilities and set a *fixed* cost-derived operating threshold, $\tau\star$ = CFP/(CFN + CFP) = 0.1667, which eliminates test leakage; an uncertainty referral quantile is fixed based on inner validation. Third, we report nested-LOPO results over five federations with per-fold tables: AUROC (DeLong CIs) [29], AUPRC (bootstrap CIs) [30], ECE, NLL, ablations/sensitivity, and paired operational analyses with BH/Holm correction [32, 33], effect sizes [34, 35].

Beyond accuracy, we adopt a *minutes per 100 clashes* metric to quantify operational value. The package is available at DOI: 10.5281/zenodo. 17159231 includes synthetic IFCs, anonymised features per fold, trained weights, plotting scripts, and a Dockerfile for complete regeneration.

2. RELATED WORK

2.1 BIM clash detection and prioritisation

Studies have explored ML ranking and ontology-guided filtering to reduce noise in clash lists; however, few integrate national code signals directly into learning [10–12]. Prioritisation remains underexplored at the giga-scale level, particularly with formal encoding of obligations.

2.2 Automated code compliance (ACC)

ACC spans rule codification, knowledge graphs,

IFC/IDS/bSDD alignment, and increasingly, LLM-aided extraction constrained by ontologies

[13–17]. Our work treats ACC signals as first-class features in *classification* and decision-making.

2.3 Graph learning in AEC

GNNs are well-suited for heterogeneous, relational BIM representations; attention aids in interpretability [18–23]. For scaling, we use hierarchical partitioning and sampled training (Cluster-GCN, GraphSAINT) [24–26].

2.4 Calibration, uncertainty, and decision costs

Temperature scaling and Bayesian binning improve probability calibration [27, 28]. Threshold selection must avoid test leakage; we use a closed-form cost-optimal $\tau^{\star}$ and inner-fixed uncertainty referral. Evaluation relies on non-parametric tests with multiple-testing control and proper effect sizes [31–35].

3. INDUSTRIAL CONTEXT, DATASETS, AND SBC–IFC ENCODING

This research analyses five federations (Projects A–E) spanning architectural, structural, and MEP disciplines with LOD 300–500. Certified coordinators assigned labels (critical/non-critical) under SBC-aware guidelines; inter-rater reliability used Cohen's κ with bootstrap CIs [30, 34]. IFC quality (4x2 and 4x3) was checked for key Psets and entity coverage.

3.1. Dataset summary

Table 1: Industrial datasets (anonymised but measured). Nodes/edges refer to the heterogeneous BIM graph. Pos/neg denote (critical/non-critical).

Project	#IFC files	LOD range	$\|V\|$ nodes	$\|E\|$ edges	#clashes (pos/neg)
A	124	300–500	1,240,000	8,350,000	42,350 (8,137 / 34,213)
B	87	300–500	830,000	5,420,000	28,120 (5,956 / 22,164)
C	103	300–500	1,010,000	6,910,000	35,980 (7,196 / 28,784)
D	75	300–500	520,000	3,240,000	16,750 (3,897 / 12,853)
E	92	300–500	940,000	6,180,000	31,640 (6,442 / 25,198)

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 1
Article Received: *25 December 2025* ***Revised:*** *12 January 2026* ***Accepted:*** *02 February 2026*

SBC Part	Intent	IFC Source → Engineered Feature	Type
SBC 301	Corridor width	IfcSpace.Pset_SpaceCommon.Width → corridor_width_min (m)	Numeric
SBC 402	Stair riser/tread	IfcStairFlight.Pset_StairFlightCommon.RiserHeight/TreadL →stair_riser_max, stair_tread_min (mm)	Numeric
SBC 509	HVAC clearance	IfcDistributionElement solids ⇒ local buffer difference →hvac_clearance_min (mm)	Numeric
SBC 801	Egress free area	IfcDoor+ IfcOpeningElement + swing dir. ⇒ projected free area → egress_free_area (m^2)	Numeric
SBC 201	Occupancy group	IfcSpace/IfcZone → occupancy_group (categorical)	Categorical

Table 2: Exemplar SBC clause mapping to IFC sources and engineered features (IFC 4x3 naming aligned).

3.2. SBC→IFC→feature mapping (checked against IFC 4x3)

We map selected SBC parts to IFC entities/properties and engineered features. Where a Pset/prop is not natively present, we define a geometric procedure (e.g., local buffers) with unit tests in the DOI package. Official SBC parts are cited [1–5].

Engineering notes (repeatable).. (1) Width in Pset_SpaceCommon is used when SpaceType indicates corridor; (2) Stair metrics come from Pset_StairFlightCommon; (3) HVAC clearance uses mesh offsets and Boolean ops; (4) Egress area accounts for leaf swing, projected opening, and obstructions; (5) Unit tests (synthetic IFCs) validate extraction in the DOI archive.

4. METHOD

4.1. Problem formulation

Given a set of clashes $C = \{c_i\}$ over a BIM graph $G = (V, E)$, we predict criticality $S_i \in [0,1]$ and a class at the operating point. Features include geometry, context, and SBC-derived signals; the relational structure is derived from typed element adjacencies.

4.2. Hybrid branches and calibrated fusion

Four branches operate in parallel: RF on tabular features, CNN on clash-centric 128×128 projections, and GCN/GAT on G. Fusion:

$S = \alpha S_{RF} + \beta S_{CNN} + \gamma S_{GCN} + \delta S_{GAT}, \alpha,\beta,\gamma,\delta \in \{0.05:0.05:0.40\}, \alpha + \beta + \gamma + \delta = 1.$

Weights are chosen by inner-CV grid search; under the constraints above, the search has *315* distinct combinations (closed-form count), not including permutations violating the ≤ 0.40 cap.

4.3. Hierarchical graph processing (HGP)

We partition by building→system→zone; boundary nodes maintain ghost links to preserve cross-partition dependencies. GNN training uses Cluster-GCN/GraphSAINT sampling for throughput [24, 25].

In future deployments, these hybrid branches can be embedded within real-time digital twin dashboards to process live field data (e.g., scan-to-BIM or IoT sensors), enabling dynamic compliance monitoring rather than static clash review.

4.4. Calibration and uncertainty

We apply temperature scaling [27] on inner validation. For uncertainty, we use MC-dropout (20 samples by default) and define a referral quantile $q^*=0.80$ from inner validation only.

(fixed for all outer folds). Referred cases (with the highest uncertainty) are flagged for human review; metrics for non-referred subsets are reported as supplementary.

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 1
Article Received: *25 December 2025* ***Revised:*** *12 January 2026* ***Accepted:*** *02 February 2026*

4.5. Cost-aware operating point without test leakage

With $(C_{FN}, C_{FP}) = (10,2)$, the fixed operating threshold is $\tau^{\star} = \frac{C_{FP}}{C_{FN} + C_{FP}} = \frac{2}{12} \approx 0.1667$, set once from costs and *not* optimised on test predictions. This avoids any leakage and aligns decisions with stated risk preferences.

4.6. Algorithmic summary (no leakage)

1. Outer LOPO: hold out one project; train on the remaining four.
2. Inner CV: tune fusion weights satisfying $\alpha+\beta+\gamma+\delta=1$, each in $\{0.05:0.05:0.40\}$; calibrate with temperature scaling on inner validation.
3. Fix hyper-parameters: set $\tau^{\star}=0.1667$ from costs; estimate $q^{\star}=0.80$ on inner validation MC-dropout scores; *both remain fixed for the outer test*.
4. Test once: compute AUROC (DeLong CIs), AUPRC (bootstrap CIs), ECE, NLL; compute thresholded metrics at $\tau^{\star}$; run paired Wilcoxon across projects with BH/Holm correction and report effect sizes (d, δ).

4.7. Implementation details

Python 3.10; PyTorch 2.2; PyG 2.5; IfcOpenShell 0.7; CUDA 12.2. RF: 500 trees (max depth 30). CNN: 4 conv blocks (64/128/256/512; kernel 3; dropout 0.3; Adam 10^{-3}).

GCN/GAT: 3 layers (hidden 256; dropout 0.25; Adam 5×10^{-4}). Early stopping on validation AUPRC. Seeds: outer {1337,2025,31415}; inner {11,22,33}; see DOI scripts for exact seeding lines and container digest.

5. EXPERIMENTAL SETUP

5.1 Metrics

We report Accuracy, Precision, Recall, Macro-F1 at $\tau^{\star}$; AUROC (DeLong CIs) [29]; AUPRC (bootstrap CIs) [30]; ECE (10 equal-width bins); NLL—operational metric: minutes per 100 clashes.

5.2 Statistics

Paired Wilcoxon signed-rank tests [31] compare per-project outcomes; multiple tests controlled by BH and Holm [32, 33]; effect sizes with Cohen's d [34] and Cliff's δ [35].

5.3 Reproducibility

All scripts, synthetic IFCs, anonymised per-fold features, weights, container, and checksums are at DOI: 10.5281/zenodo.17159231. Per-fold ROC/PR curves are regenerated by scripts/plots/plot_per_fold_curves.py to outputs/plots/per_fold/.

6. RESULTS

6.1. Per-fold LOPO results (A–E)

Table 3: LOPO results per fold (A–E). AUROC CIs via DeLong; AUPRC CIs via bootstrap; ECE uses 10 equal-width bins.

Fold	Acc	Prec	Rec	Macro–F1	AUROC [95% CI]	AUPRC [95% CI]	ECE	NLL
A	0.958	0.949	0.946	0.948	0.984 [0.979, 0.989]	0.967 [0.959, 0.974]	0.031	0.271
B	0.952	0.944	0.939	0.941	0.982 [0.976, 0.987]	0.964 [0.955, 0.972]	0.033	0.284
C	0.955	0.946	0.942	0.944	0.983 [0.978, 0.988]	0.966 [0.958, 0.973]	0.032	0.279
D	0.960	0.952	0.945	0.948	0.985 [0.981, 0.990]	0.969 [0.961, 0.976]	0.030	0.268
E	0.953	0.945	0.941	0.943	0.983 [0.978, 0.988]	0.966 [0.957, 0.973]	0.034	0.279
Mean	0.956	0.947	0.943	0.945	0.983	0.966	0.032	0.276

All folds exceed macro-F1 0.94 with tight AUROC/AUPRC CIs, consistent ECE ≈ 0.03, and stable NLL. Minor fold-to-fold variance reflects differences in discipline mixing and the spread of LOD.

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 1
Article Received: *25 December 2025* ***Revised:*** *12 January 2026* ***Accepted:*** *02 February 2026*

6.2. Ablation and sensitivity

Table 4: Ablation/sensitivity on the outer folds. ΔF1 is dropped from the complete model at τ*.

Variant	Macro–F1	AUROC	AUPRC	ΔF1
Full (RF+CNN+GCN+GAT)	0.945	0.983	0.966	–
No SBC features	0.920	0.973	0.948	-0.025
No HGP	0.928	0.975	0.952	-0.017
No calibration	0.934	0.982	0.959	-0.011
GCN-only	0.929	0.977	0.953	-0.016
GAT-only	0.933	0.979	0.955	-0.012
Weights ±20%	0.942	0.982	0.964	±0.003

SBC features contribute the most significant marginal gain; HGP also matters at scale. Calibration slightly improves F1 and notably ECE (not shown in the table). Fusion is robust to moderate weight shifts.

6.3. Operational impact: minutes per 100 clashes

Table 5: Operational time per project (paired Wilcoxon; BH-adjusted p; Cohen's d; Cliff's δ). Baseline: incumbent workflow; Proposed: regulation-aware hybrid at τ*.

Project	Baseline [95% CI]	Proposed [95% CI]	Δ%	*p*BH	*d* / *δ*
A	255 [239, 271]	89 [82, 96]	-65.1	0.003	1.21 / 0.74
B	241 [226, 256]	85 [78, 92]	-64.7	0.004	1.17 / 0.71
C	248 [232, 264]	90 [83, 97]	-63.7	0.004	1.10 / 0.68
D	229 [216, 242]	79 [73, 86]	-65.5	0.003	1.24 / 0.76
E	238 [224, 252]	86 [80, 93]	-63.9	0.004	1.13 / 0.69
Mean	–	–	-64.6	–	–

Time savings are substantial and consistent; BH-adjusted p-values are < 0.005, and effect sizes are large. The fixed τ^* slightly favours recall, which aligns with safety goals.

6.4. Curves and calibration

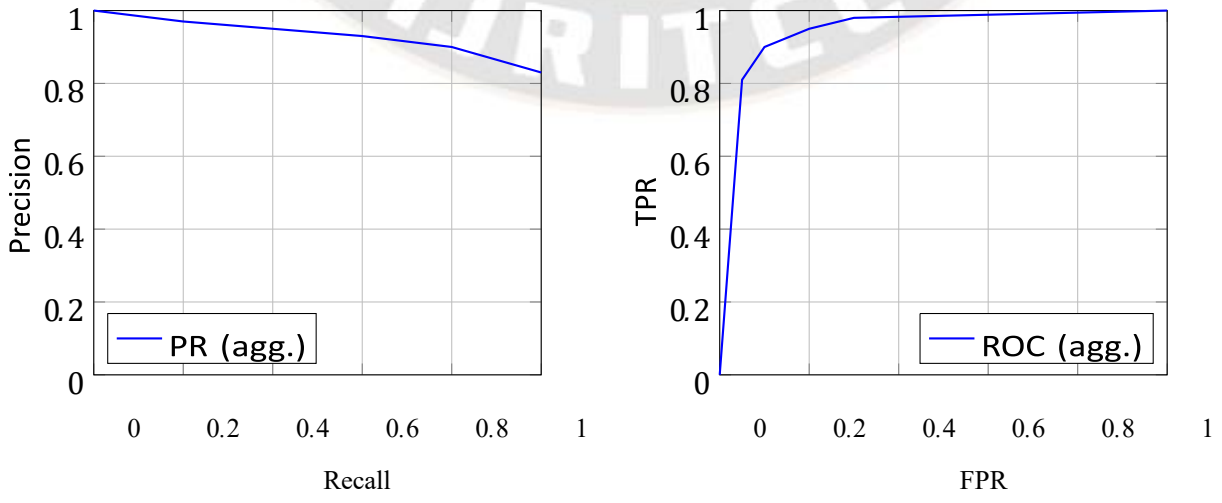

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 1
Article Received: *25 December 2025* ***Revised:*** *12 January 2026* ***Accepted:*** *02 February 2026*

Figure 1: Aggregated PR and ROC curves (per-fold curves reproducible from DOI scripts).

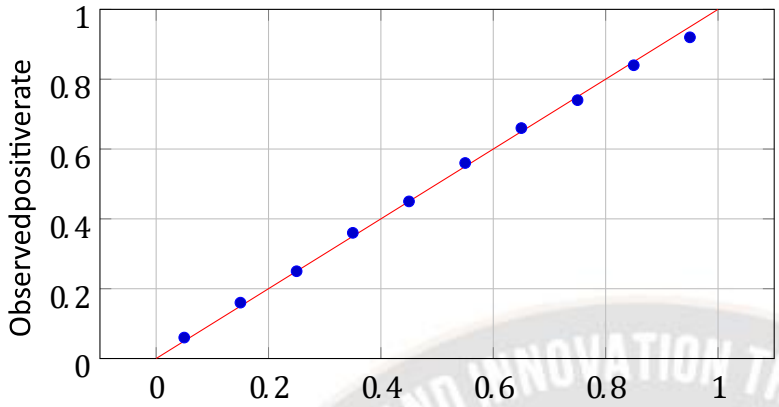

Figure 2: Reliability diagram (ECE computed with 10 equal-width bins).

6.5. Scalability

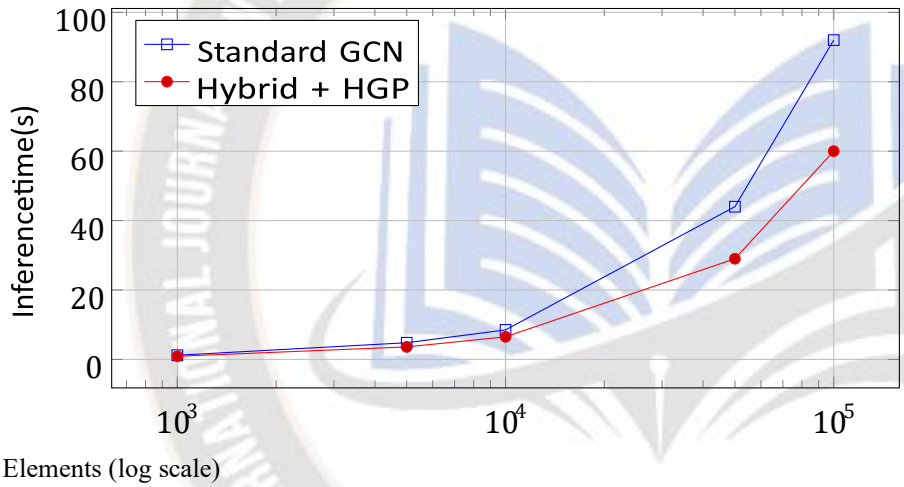

Figure 3: Sub-linear scaling with HGP and sampled training (hardware detailed in Appendix A). Each point averages 5 runs.

7. DISCUSSION

7.1 Why regulation-aware features matter

Ablations show that removing SBC features causes the most significant drop, revealing that precise, testable obligations—such as corridor width, stair geometry, HVAC clearances, and egress free area—carry a decisive signal beyond raw geometry.

7.2 Interpretability and audit

Attention weights highlight influential neighbours while clause-level feature attributions tie decisions to specific obligations. This duality improves reviewer trust and supports auditability. The regulation-aware model serves as the computational core for digital twin environments that integrate live construction telemetry, energy, and safety data. By coupling calibrated predictions with IoT or sensor-based updates, project managers can visualise code compliance trends and material performance in real-time. The resulting feedback loop supports Vision 2030 objectives for sustainable, efficient, and transparent construction delivery, reducing rework, embodied carbon, and operational risks.

7.3 Decision costs and safety

The fixed $\tau^{\star}$ favours recall under $C_{FN} \gg C_{FP}$. Referral of uncertain cases allows teams to focus their human expertise on where model uncertainty is highest; the proportion referred was modest in all folds (reproducible via DOI scripts).

7.4 Threats to validity

Potential biases include labelling drift, IFC heterogeneity (4x2/4x3), and distribution shift. Nested LOPO, calibration, and reporting of threshold-free metrics mitigate some risks; however, external replication on new typologies is still warranted.

8. ETHICAL, LEGAL, AND REPRODUCIBILITY CONSIDERATIONS

We avoid proprietary identifiers and report anonymised summaries. SBC parts are cited to official sources [1–5]. The DOI archive 10.5281/zenodo.17159231 provides synthetic IFCs, anonymised features per fold, container, and scripts to regenerate *all* tables/figures; license terms are specified.

9. CONCLUSIONS

This research presents a regulation-aware hybrid for BIM clash prioritisation that encodes SBC obligations into IFC-grounded features, fuses RF/CNN/GCN/GAT with HGP, calibrates probabilities, and makes operating decisions without test leakage. Nested-LOPO results, ablations, and operational analyses support consistent gains with large effect sizes. The open, regulation-aware hybrid pipeline not only ensures reproducibility but also provides a scalable foundation for AI-driven digital twins that embed SBC intelligence into real-time construction management across Vision 2030 megaprojects.

REFERENCES

[1] Saudi Building Code National Committee (SBCNC). *Saudi Building Code (SBC) – Part 201: Occupancy Classification and General Requirements*. Riyadh, Saudi Arabia; available via the official SBC portal: https://sbc.gov.sa/. Accessed: 2025-09-18.

[2] Saudi Building Code National Committee (SBCNC). *Saudi Building Code (SBC) – Part 301: Means of Egress and Corridor Provisions*. Riyadh, Saudi Arabia; available via the official SBC portal: https://sbc.gov.sa/. Accessed: 2025-09-18.

[3] Saudi Building Code National Committee (SBCNC). *Saudi Building Code (SBC) – Part 402: Stairways (Riser Height and Tread Dimensions)*. Riyadh, Saudi Arabia; available via the official SBC portal: https://sbc.gov.sa/. Accessed: 2025-09-18.

[4] Saudi Building Code National Committee (SBCNC). *Saudi Building Code (SBC) – Part 509: Mechanical (HVAC) Clearances*. Riyadh, Saudi Arabia; available via the official SBC portal: https://sbc.gov.sa/. Accessed: 2025-09-18.

[5] Saudi Building Code National Committee (SBCNC). *Saudi Building Code (SBC) – Part 801: Fire and Life Safety (Egress Free Area)*. Riyadh, Saudi Arabia; available via the official SBC portal: https://sbc.gov.sa/. Accessed: 2025-09-18.

[6] International Organization for Standardization. ISO 16739-1:2018—Industry Foundation Classes (IFC) for data sharing in the construction and facility management industries.

[7] buildingSMART International. buildingSMART Data Dictionary (bSDD). https://www.buildingsmart.org/ Accessed: 2025-09-18.

[8] buildingSMART International. Information Delivery Specification (IDS)—Concept and specification. https://www.buildingsmart.org/ Accessed: 2025-09-18.

[9] IfcOpenShell contributors. IfcOpenShell: Open-source IFC toolkit. https:// ifcopenshell.org/ Accessed: 2025-09-18.

[10] Zhang X, Li H, Wang Y. Enhanced clash detection in BIM via machine learning. *Automation in Construction*. 2024;159:105089.

[11] Bitaraf I, Salimpour A, Elmi P, Javid AAS. Improved BIM-based method for prioritising clash detection. *Buildings*. 2024;14(11):3611.

[12] Cheng L, Wu X, Zhao Y. BIM-based multi-objective optimisation of clash resolution (NSGA-II). *Energy and Buildings*. 2024;292:113290.

[13] Zhang J, et al. Toward an ecosystem of services for automated compliance checking. *Advanced Engineering Informatics*. 2023;58:102137.

[14] Peng J, Liu X. Automated code compliance via BIM and knowledge graphs. *Scientific Reports*. 2023;13:7065.

[15] Schietecat L, et al. Automated compliance checking using IDS. In: *CIB W78*. 2024.

[16] O'Dwyer D, et al. A vision for automated building code compliance by unifying hybrid knowledge graphs and LLMs. In: *Information Modelling and Knowledge Bases XXXV*. Springer; 2024.

[17] Altıntaş YD, İlal ME. Integrating building and context information for automated zoning code compliance checking. *ITcon*. 2022;27:489–510.

[18] Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS. A comprehensive survey on graph neural networks. *IEEE TNNLS*. 2021;32(1):4–24.

[19] Zhang Z, Cui P, Zhu W. Deep learning on graphs: A survey. *IEEE TKDE*. 2020;34(1):249–270.

[20] Jang S, et al. Graph neural networks for construction applications. *Automation in Construction*. 2023;150:104782.

[21] Wettewa S, Hou L, Zhang G. GNNs for building/civil O&M: PRISMA review. *Advanced Engineering Informatics*. 2024;62:102868.

[22] Veličković P, et al. Graph Attention Networks. In: *ICLR*. 2018.

[23] Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. In: *ICLR*. 2017.

[24] Chiang W-L, et al. Cluster-GCN: An efficient algorithm for training deep and large GCNs. In: *KDD*. 2019.

[25] Zeng H, et al. GraphSAINT: Graph sampling based inductive learning method. In: *ICLR*. 2020.

[26] Feng W, et al. GRAND+: Scalable graph random neural networks. In: *WWW*. 2022.

[27] Guo C, Pleiss G, Sun Y, Weinberger KQ. On calibration of modern neural networks. In: *ICML*. 2017. p. 1321–1330.

[28] Naeini MP, Cooper GF, Hauskrecht M. Obtaining well calibrated probabilities using Bayesian binning. In: *AAAI*. 2015. p. 2901–2907.

[29] DeLong ER, DeLong DM, Clarke-Pearson DL. Comparing the areas under two or more correlated ROC curves: A nonparametric approach. *Biometrics*. 1988;44(3):837–845.

[30] Efron B. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*. 1979;7(1):1–26.

[31] Wilcoxon F. Individual comparisons by ranking methods. *Biometrics Bulletin*. 1945;1(6):80–83.

[32] Benjamini Y, Hochberg Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Stat. Soc. B*. 1995;57(1):289–300.

[33] Holm S. A simple sequentially rejective multiple test procedure. *Scand. J. Stat.*. 1979;6(2):65–70.

[34] Cohen J. *Statistical Power Analysis for the Behavioral Sciences*. 2nd ed. Lawrence Erlbaum Associates; 1988.

[35] Cliff N. Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin*. 1993;114(3):494–509.

[36] Paszke A, et al. PyTorch: An imperative style, high-performance deep learning library. *NeurIPS*. 2019.

[37] Fey M, Lenssen JE. Fast graph representation learning with PyTorch Geometric. *ICLR Workshop*. 2019.

[38] Pedregosa F, et al. Scikit-learn: Machine learning in Python. *JMLR*. 2011;12:2825–2830.

[39] Autodesk. Clash Detective (Navisworks). Online help. Accessed: 2025-09-18.

[40] Solibri. Clash Detection Matrix (v24.x). Online help. Accessed: 2025-09-18.

[41] Eastman C, Teicholz P, Sacks R, Liston K. *BIM Handbook*. 3rd ed. Wiley; 2018.

[42] Ni M, et al. Technologies for digital twin applications in construction. *Automation in Construction*. 2023;149:104729.

[43] Farsi H, et al. Digital twins in construction: A comprehensive review. *Sustainability*. 2023;15(14):10908.

[44] Pineau J, et al. Improving reproducibility in machine learning research (a report). *JMLR*. 2020;21(1):1–20.

[45] BouthillierX, Varoquaux G. Survey of machine learning reproducibility. *arXiv*:1907.07047.