A BIM and LLM Framework for Automated Construction and Demolition Waste Management

Nolan Porther[1], Amirhossein Nourbakhshrezaei[1] and Mojgan Jadidi[1, *]

[1] Lassonde School of Engineering, York University, Toronto, Canada - {nolanpo; amirnbr2; mjadidi}@yorku.ca

Keywords: Construction and Demolition Waste (C&DW); Building Information Modelling (BIM); Large Language Models (LLM); Retrieval-Augmented Generation (RAG); Bayesian Cost Estimation; Circular Economy

Abstract

Artificial Intelligence (AI) integration has become an essential of modern AEC workflows, yet it has failed to gain a position in waste management. This gap is particularly prominent given the urgent environmental and legal imperatives for the sector to mitigate its demolition outputs. Existing approaches to waste classification and diversion cost estimation rely on manual interpretation of project documentation, a process that is both resource-intensive and structurally incompatible with the machine-readable data environments established by Building Information Modelling (BIM). This paper presents a framework that bridges Industry Foundation Class (IFC) compliant BIM data and Large Language Model (LLM) capabilities to automate Construction and Demolition Waste (C&DW) classification and probabilistic cost optimisation. The framework utilizes IfcOpenShell to extract element geometry and material data, channeling this information into a Retrieval-Augmented Generation (RAG) pipeline. To ensure rigorous compliance during classification, a FAISS-indexed knowledge base grounds a locally deployed Llama3 model against the specific mandates of Province of Ontario, Canada regulation 102/94. Diversion cost scenarios are computed through a Bayesian cost module coupled to a multi-objective genetic algorithm (MOGA) optimiser. Th proposed approach is evaluated against a labelled dataset of 104 IFC type-and-material combinations, the RAG classifier. Performance thresholds were established a piori based on multi-class classification benchmarks and Bayesian cost model uncertainty tolerances. The framework achieved a macro-average F1 of 0.84 and overall accuracy of 88%, satisfying the minimum criteria for automated C&DW characterization under Ontario Regulation 102/94.

1. Introduction

Construction and demolition waste (C&DW) is, by most measures, the largest single solid waste category in the Ontario, Canada built environment. More than a quarter of all material deposited in provincial landfills annually is C&DW (Ministry of the Environment, Conservation and Parks, 2023), a proportion that has remained obstinately stable despite two decades of sustained regulatory effort. The formal pre-project waste audit obligation under Ontario Regulation 102/94 has been on the books since 1994. The practice on the ground has changed very little: assessors still walk sites, cross-reference paper drawings and apply professional judgment to categorize materials in a process that is slow, variable across practitioners and wholly disconnected from the digital project data that modern procurement workflows routinely produce (Akinade et al, 2017). That gap between regulatory intent and the tools available to fulfill it is the central problem motivating this work.

The adoption of Building Information Modelling (BIM) as a project delivery standard has, in principle, already resolved the data availability problem. Industry Foundation Class (IFC) compliant models encode element geometry, material composition, and classification hierarchies in a standardized machine-readable schema that can be parsed programmatically. The growing mandate for ISO 19650-aligned information management on public infrastructure projects means this data is increasingly available precisely at the point where waste audit obligations arise. What has been missing is a credible pipeline for translating raw model data into regulatory-compliant waste classifications and honest cost estimates that handles the inevitable gaps in BIM authoring quality and reconciles model data with unstructured project specifications.

Recent advances in Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) offer a promising path towards closing that gap. LLMs perform well on domain-specific classification tasks when supplied with the right contextual grounding, and RAG substantially improves factual alignment by retrieving relevant regulatory passages at inference time rather than relying on pre-trained knowledge alone (Lewis et al., 2020). The critical challenge is not the classification capability of these tools, but whether integrated pipelines built around them can achieve the levels of accuracy, traceability and regulatory alignment required for pre-demolition waste audits within specific provincial jurisdictions. This gap remains insufficiently addressed in the current literature.

The proposed framework integrates IFC model parsing via IfcOpenShell, a FAISS-indexed RAG knowledge base from Province of Ontario, Canada regulation 102/94 and 103/94, a locally deployed Llama3 LLM with deterministic override rules. Then a Bayesian cost optimization model coupled applying Multi-Objective Genetic Algorithm (MOGA) optimizer.

The remainder of the paper is organized as follows: Section 2 situates the contribution in the existing literature. Section 3 describes the methodology in full. Section 4 presents the case study and results. Section 5 states limitations and future directions for continuing study. Section 6 concludes the paper by highlighting the contribution, obtained results and future direction.

2. Background

2.1 BIM-Based Waste Estimation

BIM-based approaches to waste estimation have attracted growing research attention over the past decade, motivated largely by the demonstrated inadequacy of conventional generation-rate ratios for project-specific quantity prediction. Akinade et al. (2018) showed that direct material take-off from BIM geometry can substantially outperform ratio-based benchmarks in deconstruction planning and Won and Cheng (2017) provided a systematic review of BIM opportunities across the C&DW lifecycle. The IFC schema's IfcElement hierarchy provides a principle basis for programmatic extraction, where IfcQuantitySet entities carry pre-computed volumes. The Open CASCADE geometry kernel underlying IfcOpenShell provides a fallback computation path for elements lacking explicit quantity assignments (Verheij and Lipman, 2011). However, the current states of art have not yet adequately addressed the downstream classification problem. Extracting quantities from a model is considerably easier than mapping those quantities onto a regulatory waste taxonomy for a specific jurisdiction, and most published frameworks either use generic international schemes or stop short of the classification step entirely.

2.2 LLMs and Retrieval-Augmented Generation (RAG) in Construction

The application of LLMs to construction document understanding is a more recent development, but early results are encouraging. Zheng et al. (2024) demonstrated prompt-engineered LLM performance on construction specification tasks, and growing work in adjacent domains suggests that retrieval augmentation substantially improves classification reliability for regulatory and standards-based tasks where factual grounding matters (Gao et al., 2023). The key mechanism is straightforward. Rather than relying on whatever relevant knowledge the model absorbed during pre-training, RAG supplies retrieved passages from a curated external knowledge base at inference time. This makes the classification decision verifiable and traceable to specific source documents (Lewis et al., 2020). For waste classification against a provincial regulatory taxonomy, this is not merely a nice-to-have; it is a prerequisite for any credible claim to audit applicability.

2.3 Probabilistic Cost Modelling and Multi-Objective Optimization

Multi-objective optimization of construction waste logistics has a reasonably well-developed literature, with Multi-Objective Genetic Algorithm (MOGA)-based approaches demonstrating good performance on problems requiring simultaneous optimization of cost, carbon, and diversion objectives (Tam and Tam, 2008). The constant drawback is that cost parameters are treated as deterministic quantities. This is a simplification that does not reflect the volatility characterizing real waste markets. Regional tipping fees fluctuate, recovered material revenues track commodity prices, and haul cost estimates depend on facility availability that can change on short notice (Waste Diversion Ontario, 2014). Bayesian approaches to cost uncertainty in construction are well-established, but their integration with multi-objective waste optimization pipelines has received little attention (Sun et al, 2014). That integration is one of the design contributions of the present work.

2.4 Research Gap

Prior work has developed BIM-based methods for extracting material and element-level information (Akinade et al., 2017; Won and Cheng, 2017) and established regulatory classification systems such as CSA Z795 (CSA Group, 2019), while recent advances in LLMs enable automated classification from semi-structured data (Gao et al., 2023; Lewis et al., 2020). However, these approaches are typically applied in isolation. BIM workflows rely on manual or rule-based classification, and LLM-based methods are rarely grounded in structured BIM semantics or regulatory standards. As a result, an integrated framework that unifies BIM-based semantic extraction, LLM-driven regulatory classification, and standards-compliant waste quantification remains absent. The main contribution here is integration and evaluation of such data against a concrete performance standard in a regional context, in our study Ontario, Canada.

3. Methodology

3.1 System Architecture and Digital Twins

The proposed framework is organized as a four-stage sequential pipeline, summarized in Table 1. The system architecture is designed with operational digital twins’ deployment as a hypothesis. The database schema, query interface and optimization module are all built to eventually accept live sensors and Internet of Things (IoT) data streams. The current implementation operates on static offline project data and should be understood as a descriptive-to-informative twin in the terminology of Grieves and Vicker (2017). References to "real-time" interaction in this paper refer to the natural-language query interface, not to live data ingestion from construction sites. This distinction is important because it influences how the results are interpreted and what can realistically be claimed about operational readiness.

Stage	Process	Output
4A	IFC parsed via IfcOpenShell; geometry and material attributes extracted	DataFrame: IFC_Type, Material, Volume, Mass
4B	RAG knowledge base from Ontario Reg. 102/94 and 103/94; FAISS index built	Regulatory context retrieval system
4C	Llama3 LLM invoked per element with override rules and retrieved regulatory context	Classified DataFrame: waste_category, disposal_route, confidence, reasoning
4D	Unique IFC/material combinations manually labeled against Reg. 102/94; evaluated with scikit-learn	Classification report; macro-average F1 ≥ 0.80

Table 1. Summary of the automated BIM-to-waste classification pipeline (Stages 4A–4D).

The four stages proceed as follows. Stage 4A parses an IFC-format BIM model to extract element geometries, material attributes and quantity data. Stage 4B constructs a regulatory knowledge base from Ontario Regulation 102/94 and 103/94 and indexes it uses FAISS for fast similarity retrieval. Stage 4C deploys a locally hosted Llama3 LLM with a structured classification prompt, retrieved regulatory context and deterministic override rules to assign waste categories and confidence scores to each building element. Stage 4D exports unique IFC type-and-material combinations for manual labeling, merges labeled data with classifier outputs and evaluates performance using scikit-learn classification metrics against a pre-specified macro-average F1 threshold. In the final steps of Stage 4D, a Bayesian cost module and MOGA optimizer operate on classified waste streams to generate probabilistic cost estimates and Pareto-optimal diversion scenarios. The overall pipeline is illustrated in Figure 1.

3.2 IFC Model Parsing and Attribute Extraction (Stage 4A)

IFC models are ingested using IfcOpenShell, which provides a Python-native interface to the IFC schema's entity hierarchy. All IfcBuildingElement subclasses relevant to C&DW generation are extracted which encompasses walls, slabs, beams, columns, doors, windows, stairs and ceiling coverings. For each element, four attributes are retrieved: IFC entity type, Revit family name, object type designation, and material assignment. Geometric quantities are taken from IfcElementQuantity property sets where available; where they are absent, volume is computed from element shape representations using the Open CASCADE geometry kernel. Material densities are assigned from a lookup table keyed to waste category, drawing on characterization data from Cochran et al. (2007) and Waste Diversion Ontario (2014), enabling downstream mass estimation in metric tonnes. Elements for which neither material assignment nor object type provides a usable classification signal are flagged with an 'Unknown' material attribute and passed to the LLM in Stage 4C for inference-based resolution.

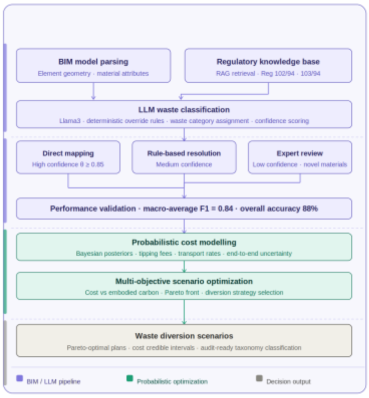

Figure 1. Four stages automated C&DW waste management framework from BIM model parsing through waste diversion scenario output.

3.3 Regulatory Knowledge Base and RAG Infrastructure (Stage 4B)

The regulatory knowledge base is built from the full text of Province of Ontario, Canada regulation 102/94 and its companion Regulation 103/94, segmented into overlapping semantic chunks of approximately 300 tokens with a 50-token boundary overlap following the chunking strategy described by Siriwardhana et al. (2023). Chunk embeddings are computed using a sentence-transformer model and indexed in a FAISS flat inner-product index (Johnson et al., 2019). At inference time, a retrieval query is assembled for each building element by concatenating its IFC entity type, family name, and object type designation, and the top 3 most semantically similar regulatory passages are retrieved and prepended to the LLM system prompt. The classifier’s output is uniquely grounded in local requirements, setting it apart from generic international models. Specifically, waste category assignments are mapped to Ontario’s regulatory framework rather than global standards like the European Waste Catalogue. The traceability this design affords is an important property for any system intended to support regulatory compliance.

3.4 LLM-Based Classification with Override Rules (Stage 4C)

Waste classification is performed by Llama3, an 8-billion parameter open-weight model from Meta AI, served locally through the Ollama inference runtime. Running the system locally rather than through a cloud API was driven by a real constraint in this space: IFC models often include sensitive design details and cost information that, under typical Canadian construction contracts, cannot be transmitted to external services (Succar, 2009). The system prompt clearly defines the eight categories from Ontario Regulation 102/94 and requires the model to return results in a structured JSON format, including waste_category, disposal_route, confidence, and reasoning.

As testing progressed, fourteen deterministic override rules were embedded directly in the prompt to address recurring misclassification patterns observed during evaluation. Most critically, the model's tendency to assign IfcWindow elements without explicit material data to Mixed/Other rather than Glass, and its misclassification of partition walls as Concrete/Masonry rather than Drywall/Gypsum. Taxonomy mapping from LLM outputs to CSA Z795 and EPA material codes follows a three-tier confidence protocol. Tier 1 covers high-confidence assignments where the model's posterior classification probability meets or exceeds $\theta \geq 0.85$; these are mapped directly without further review. Tier 2 handles medium-confidence outputs through a rule-based disambiguation layer applying material composition heuristics and cross-referencing the element's IFC property sets. Tier 3 flags low-confidence or materially novel descriptions which includes composite assemblies, proprietary finish systems and elements with completely absent material data for human-in-the-loop review. Every classification decision is logged with provenance metadata including the retrieved regulatory passage, confidence score and resolution tier, creating an auditable record supporting both regulatory reporting and future model fine-tuning.

3.5 Bayesian Cost Modelling and MOGA Coupling (4D)

Unit cost parameters such as disposal tipping fees, transportation rates and recovered material revenues are treated as uncertain quantities and modelled as probability distributions rather than point estimates. Prior distributions are specified from regional Ontario waste market data and updated via Bayesian inference against available project-specific cost records, so that higher-quality project data narrow the posteriors while projects with limited cost history carry wider uncertainty. Rather than reducing probability distributions to single expected values prior to optimization, the MOGA samples directly from the distributions when evaluating objective values for each diversion scenario. As a result, the objective values vary based on the underlying cost uncertainty, and the Pareto front becomes a range of possible non-dominated solutions rather than a single deterministic curve.

3.6 End-to-End Uncertainty Propagation

The framework treats uncertainty as a first-class quantity tracked from the moment of extraction through to the final scenario recommendation. At the extraction stage, LLM classification confidence scores and IFC geometric tolerance metadata are both passed forward as informative signals. These feed into Bayesian prior so that the posterior distributions over unit costs reflect not just market volatility but the quality of the classification evidence that produced them. When the MOGA presents a scenario to the decision-maker, the associated cost range is a 90th percentile posterior credible interval that captures both sources of uncertainty: the inherent unpredictability of waste markets and the epistemic limits of the classification process.

4. Case Study and Results

4.1 Case Study Model Description

The framework was applied to an open-source IFC institutional building model sourced from the Open IFC Model Repository maintained by the University of Auckland[1]. The model represents a two-storey commercial building with a reinforced concrete and masonry structural system, glazed curtain wall facades, and a mix of interior partition, mechanical, and finish elements. Figure 2 shows the model georeferenced within its site context as visualized through the BIM-to-GIS integration layer that underpins the framework's spatial cost modelling capability. No project specifications were available from the open repository; accordingly, the RAG retrieval context relied exclusively on Ontario Regulation 102/94 and 103/94 source documents, representing a more constrained evaluation condition than a live project deployment would provide.

Figure 2. IFC building model visualized in ArcGIS Scene with site context (Open IFC Model Repository, University of Auckland).

4.2 Classification Performance

The IFC parsing module extracted 1,847 building elements across 14 element categories. Material associations were resolved directly from model data for 94 percent of elements; the remaining six percent required rule-based fallback assignment based on element type and contextual signal in family naming conventions. The RAG-augmented LLM classifier was evaluated against a manually labeled ground-truth dataset of 104 unique IFC type-and-material combinations, with expert labels assigned directly against the Ontario Regulation 102/94 category taxonomy. Six waste categories were represented in the labeled set: Drywall/Gypsum (32 combinations), Concrete/Masonry (19), Glass (19), Mixed/Other (14), Wood (12), and Metal (8). Following iterative refinement of override rules and correction of three ground-truth labeling inconsistencies identified through diagnostic analysis, the framework achieved a macro-average F1 of 0.84 and an overall accuracy of 88 percent, both exceeding the pre-specified performance thresholds. Per-class results are reported in Table 2.

Glass and Metal both achieved perfect F1 scores of 1.00 following the introduction of override rules targeting IfcWindow entities and steel stair assemblies. Drywall/Gypsum reached F1 0.98, driven by the strong override signal for partition wall family name patterns. Concrete/Masonry returned a balanced F1 of 0.95 with equal precision and recall. The two problematic categories are Mixed/Other, where high precision (0.80) paired with very low recall (0.29) indicates systematic misrouting of genuinely ambiguous elements to more specific categories and Wood, where perfect recall against low precision (0.55) reflects over-prediction for door assemblies whose material designation is absent in the source model. Both patterns trace to the same root cause: a significant proportion of door, composite floor and stair elements in this model carried no material specification, leaving the classifier without sufficient signal for deterministic assignment.

Category	Prec.	Rec.	F1	Supp.	Route
Concrete/Masonry	0.95	0.95	0.95	19	Recycling
Drywall/Gypsum	1.00	0.97	0.98	32	Landfill
Glass	1.00	1.00	1.00	19	Recycling
Metal	1.00	1.00	1.00	8	Recycling
Mixed/Other	0.80	0.29	0.42	14	Landfill
Wood	0.55	1.00	0.71	12	Recycling
Macro Average	**0.88**	**0.87**	**0.84**	**104**	—

Table 2. Per-class classification metrics for waste category assignment (n = 104 element combinations).

4.3 Cost Estimation and Optimization

Bayesian posterior distributions over unit costs were generated for all 23 waste streams identified in the classified element inventory, drawing on regional tipping fee schedules, transportation rate data, and recovered material revenue estimates from Ontario waste management records. The MOGA was configured with a population of 200, a generation count of 500, a crossover rate of 0.85 and a mutation rate of 0.02, calibrated through preliminary sensitivity analysis. The algorithm converged to a Pareto front of 47 non-dominated solutions spanning the cost-carbon trade-off space. The cost-optimal frontier solution achieved an estimated diversion rate of 78 percent, with an associated waste management cost uncertainty band of ±18 percent at the 90th percentile posterior credible interval.

That uncertainty band does not mean the model is imprecise, just honest. The ±18 percent reflects genuine volatility in Ontario 's regional waste markets and the epistemic limits of the classification evidence for some waste streams. A model returning a tighter band without better data would be projecting false confidence. The demonstrated results are automated element-level waste stream extraction and classification at 88 percent accuracy against an Ontario-specific regulatory taxonomy, and probabilistic cost estimation with credible-interval reporting. The performance benchmarks cited in the literature which says waste quantity estimation within ±5 percent of independently audited values and cost savings exceeding 20 percent relative to conventional estimation practice are proposed here as empirical validation targets for future work against projects where audit records are available. They are not claims about the present experiment.

5. Limitations and Future Work

Several limitations constrain the generalizability of the results reported here. The empirical evaluation rests on a single open-source building model. Performance will vary across building typologies particularly curtain wall towers where glazing systems introduce complex composite assemblies that stress the

[1] openifcmodel.cs.auckland.ac.nz

classification pipeline. Also, this will be seen in across jurisdictions with different waste audit taxonomies. The ground-truth labeling was conducted by a single domain expert, which introduces the possibility of inter-annotator variability that a multi-rater study would quantify. The absence of independent waste audit data for the case study building means that the quantity estimates produced by the IFC parsing module have not been validated against physical reality; the framework's end-to-end quantity accuracy claim is at this stage a design target derived from the literature, not a demonstrated result. The routing and facility allocation module referenced in the broader framework description was not implemented or tested in this study. Finally, the LLM extraction pipeline has not been stress-tested against the full range of material descriptions that might appear in Ontario construction projects, including legacy composite assemblies, proprietary finish systems, and hazardous materials triggering Regulation 347 obligations.

Future development will proceed along three principal lines. First, multi-site deployment across a range of building typologies and project sizes will provide the empirical foundation needed to substantiate the quantity accuracy and cost reduction claims currently proposed as benchmarks. Second, the routing and facility allocation module will be implemented and validated, completing the optimization pipeline and enabling the spatial cost modelling capability that the framework's GIS integration was designed to support. Third, the LLM component will be fine-tuned on a curated Ontario-specific C&D waste characterization corpus, with the goal of reducing reliance on manually specified override rules for edge cases and improving recall for the Mixed/Other and Wood categories where material specification gaps currently suppress performance.

6. CONCLUSIONS

The proposed framework in this paper cultivated from frustration that is familiar to anyone who has watched a pre-demolition waste audit get conducted with a clipboard and a set of paper drawings on a project that has a perfectly good BIM model sitting on a server somewhere. The data to do this better and more efficient has existed for years. This paper demonstrates is potential of existing BIM data in this process by introducing the pipeline that encompasses IFC parsing, RAG-augmented regulatory classification, Bayesian cost uncertainty propagation and multi-objective scenario optimization, end to end and validation. The results shows that the proposed system performs well enough to matter. A macro-average F1 of 0.84 and an overall accuracy of 88 percent against a manually labeled expert benchmark are not headline numbers pulled from a favourable experimental condition. They are the result of working through the failure modes, correcting the labeling inconsistencies and writing the override rules that the real disorder of IFC-authored models requires. That work is part of the contribution. As summary, the main contribution of this paper is to show that the architecture works, that the integrations hold, and that the approach clears the threshold required for pre-demolition waste audit applicability under Ontario Regulation 102/94. That is a meaningful place to have arrived, and a credible place from which to continue.

It is worth to mention that the developed system is not an operational digital twin, but a decision support pipeline designed to become one, and there is real distance still to cover between those two things. The quantity accuracy and cost reduction benchmarks cited in the discussion are targets drawn from the literature, not outcomes demonstrated here. The single-building case study establishes proof of concept, not generalizability, and the authors are under no illusion that one institutional model in Ontario settles the question. These limitations are claimed here will be the main avenue for future research direction. The future works include multi-site validation, routing module implementation and domain fine-tuning of the LLM component.

Acknowledgements

The authors acknowledge support from York University for providing research space, University of Aukland Open data platform, and Natural Science and Technology Research Council (NSERC) for providing funding support.

References

Akinade, O.O., Oyedele, L.O., Bilal, M., Ajayi, S.O., Owolabi, H.A., Alaka, H.A., Bello, S.A., 2017. BIM-based approaches for construction and demolition waste estimation and minimisation: A review. *Resources, Conservation and Recycling*.

Akinade, O.O., Oyedele, L.O., Ajayi, S.O., Bilal, M., Alaka, H.A., Owolabi, H.A., Bello, S.A., 2018. Waste minimisation through deconstruction: A BIM based Deconstructability Assessment Score (BIM-DAS). Resources, Conservation and Recycling, 129, 54–65.

Cochran, K.M., Townsend, T.G., Reinhart, D., Heck, H., 2007. Estimation of regional building-related C&D debris generation and composition: Case study for Florida, US. Waste Management, 27(7), 921–931.

CSA Group, 2019. Z795-18: Recovered Materials — Construction, Renovation and Demolition. Canadian Standards Association, Mississauga, ON.

Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, H., 2023. Retrieval-Augmented Generation for Large Language Models: A Survey. arXiv:2312.10997.

Grieves, M., Vicker, J., 2017. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In: Kahlen, F.J., Flumerfelt, S., Alves, A. (Eds.), Transdisciplinary Perspectives on Complex Systems. Springer, Cham, pp. 85–113.

Johnson, J., Douze, M., Jégou, H., 2019. Billion-Scale Similarity Search with GPUs. IEEE Transactions on Big Data, 7(3), 535–547.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., Kiela, D., 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Advances in Neural Information Processing Systems, 33, 9459–9474.

Meta AI, 2024. Llama 3 Model Card. Meta Platforms, Inc., Menlo Park, CA, USA.

Ministry of the Environment, Conservation and Parks, 2023. Ontario's Waste-Free Ontario Act: Framework for Resource Recovery and Circular Economy. Queen's Printer for Ontario, Toronto, ON.

Ontario Regulation 102/94, 1994. Waste Audits and Waste Reduction Work Plans. Environmental Protection Act, R.S.O. 1990, c. E.19. Government of Ontario, Toronto, ON.

Siriwardhana, S., Ali, S.S.N., Ziems, M., Shi, T., 2023. A Study of Factors Influencing the Performance of RAG-Based Retrieval Systems for Code. arXiv:2307.03987.

Succar, B., 2009. Building Information Modelling Framework: A research and delivery foundation for industry stakeholders. Automation in Construction, 18(3), 357–375.

Sun, W., An, C., Li, G., Lv, Y., 2014. Applications of inexact programming methods to waste management under uncertainty: current status and future directions. *Environmental Systems Research*, 3(1), 1–15.

Tam, V.W.Y., Tam, C.M., 2008. Waste reduction through incentives: A case study. Building Research and Information, 36(1), 37–43.

Verheij, T., Lipman, R., 2011. IfcOpenShell: A Technical Overview. In Proceedings of the 28th International Symposium on Automation and Robotics in Construction, Seoul, South Korea.

Waste Diversion Ontario, 2014. Construction, Renovation and Demolition Waste in Ontario: Characterization and Management. Waste Diversion Ontario, Toronto, ON.

Won, J., Cheng, J.C.P., 2017. Identifying potential opportunities of BIM for construction and demolition waste management and minimization. Automation in Construction, 79, 3–18.

Zheng, Z., Lu, X.Z., Chen, K.W., Zhou, Y., Lin, J.R., 2024. Prompt engineering for building construction. Advanced Engineering Informatics, 61, 102454.