

Enabling a Material Re-Use Workflow with LLM-Extracted Metadata for Circular Construction

Peter Waldert , Lukas Gschweilt , Benedikt Kantz , Matthias Lang-Raudaschl ,
Milena Stavric , and Tobias Schreck 




Figure 1: The ReUseBIM dashboard, federating reclaimed-component listings from Digital Material Banks (DMBs) for CO₂- and cost-aware material reuse planning. Photographs from [1].

Abstract—Construction accounts for 50% of European material consumption, 36% of EU waste, and 39% of global energy-related greenhouse gas emissions, making material reuse a pressing climate priority. Reclaimed components exist in abundance, but reuse remains impractical within typical project schedules: available stock is scattered across dozens of fragmented Digital Material Banks (DMBs) - web portals that list components salvaged from demolition or refurbishment - and the data on those platforms is frequently incomplete; critical attributes are buried in PDFs rather than exposed as searchable fields. Developed in close collaboration with practising architects, ReUseBIM addresses both obstacles. A dashboard aggregates listings from DMBs into a single metadata-searchable view and provides Building Information Modelling (BIM)-ready Industry Foundation Classes (IFC) files for drag-and-drop into any BIM-capable software, helping architects to efficiently search, compare, and select reclaimed components without leaving their existing workflow. Our enrichment pipeline uses an Large Language Model (LLM) to extract missing attributes from attached documents, presenting each detected value to the architect for confirmation before writing it into the IFC file. Together, these make specifying a reclaimed component as straightforward as specifying a new one.

Index terms—Material Re-Use, Circular Economy, Visual Analytics, Building Information Modeling, Sustainability

1 INTRODUCTION

Construction is the most resource-intensive industry in industrialized countries, accounting for 50% of European material consumption,

- Peter Waldert is with Institute of Visual Computing, Graz University of Technology. E-mail: peter.waldert@tugraz.at.
- Lukas Gschweilt is with Institute of Architecture Technology, Graz University of Technology.
- Benedikt Kantz is with Institute of Visual Computing, Graz University of Technology.
- Matthias Lang-Raudaschl is with Institute of Architecture Technology, Graz University of Technology.
- Milena Stavric is with Institute of Architecture and Media, Graz University of Technology.
- Tobias Schreck is with Institute of Visual Computing, Graz University of Technology.

36% of EU waste, and 39% of global energy-related greenhouse gas emissions [19]. Reusing reclaimed building components instead of manufacturing new ones is among the most direct ways to reduce that footprint, and some municipalities now mandate it [13]. Although materials are abundant, reuse is rarely the default, because discovery and specification workflows remain fragmented and inefficient [2, 6].

Architects working with reclaimed components face two compounding challenges. First, available stock is spread across numerous disconnected Digital Material Banks (DMBs) - web portals operated by salvage organisations, municipalities, or commercial platforms that list components available for reuse (also described as material market-places) [14] - leaving no centralised point for comparison. Second, the data on those platforms is frequently incomplete: critical attributes such as fire resistance class or thermal transmittance, which are required




Figure 2: Systematic Flow of the **ReUseBIM** data preparation and usage pipeline. The dashboard, the main component contributed by this work, is outlined in teal.

to produce material passports and verify regulatory compliance, are buried in attached PDFs instead of exposed as searchable fields, or absent altogether. The result is that specifying a reclaimed component requires manual searching across platforms, followed by individual modelling of each shortlisted part - an overhead that makes reuse impractical within a typical project schedule.

These challenges, and the requirements we derive from them, emerged from a co-design process with practising architects. One senior architect with more than ten years of practice, now working in academia, took part throughout, complemented by semi-structured interviews with three practising architects in different professional roles [6]. Across several workshops we used on-paper prototyping and scenario walkthroughs to translate everyday reuse tasks into user stories, which we condensed into three design requirements:

R1: Federated discovery - A single interface should aggregate component listings from multiple **DMBs** and expose them as a searchable, comprehensive view.

R2: BIM-native export - **Building Information Modelling (BIM)** geometry should be generated automatically per component, as an **Industry Foundation Classes (IFC)** file, so that architects can drag-and-drop candidates directly into their existing **BIM** software without manual modelling.

R3: Attribute completeness - Attributes absent from marketplace listings but present in attached documentation should be extractable and writeable into the **IFC** file in a form suitable for structured comparison and material passport generation.

This paper presents **ReUseBIM**, developed in close collaboration with practising architects with domain expertise in circular construction - to integrate directly into existing **BIM** workflows. It combines

(1) a dashboard that aggregates listings from **DMBs** into a single metadata-searchable view, with automatically generated **IFC** files [7] ready to drag-and-drop into any **BIM**-capable software, and

(2) our enrichment pipeline [8], which uses a **Large Language Model (LLM)** to extract missing attributes from unstructured documentation, presenting detected values to the architect for confirmation before writing them into the **IFC** file.

By addressing the principal bottlenecks in current reuse workflows, **ReUseBIM** aims to make selecting a reclaimed component as straightforward as specifying a new one, thereby translating a climate-relevant opportunity into mainstream practice. Our earlier work built the pipelines for connecting salvage centres to **BIM** software and translated interviews with practitioners into user stories [6, 7]. This paper contributes the interface layer, focusing on the interface between **DMB** data and the architect’s **BIM** tooling, with an **LLM**-based enrichment pipeline [8] and a requirements-driven design. An overview of the full pipeline is shown in Figure 2.

2 RELATED WORK

Existing tools for circular construction largely operate within design-stage environments and do not connect to **DMBs** or native **BIM** software.

RhinoCircular [9] is a **Grasshopper** plugin that scores design decisions for circularity and generates material passports, but works only within **Rhino/Grasshopper**. **Augmented Reuse** [18] uses an **Augmented Reality (AR)** measurement app to capture components from buildings due for demolition, supporting creative reuse in early design, again without **BIM** or **DMB** integration. Brütting et al. [3] optimise reusable steel structures from a shared kit of parts, addressing structural reuse instead of discovery or **BIM** integration.

On the product-data side, Digital Product Passports [5] provide a framework for standardized, machine-readable data for construction products. Passports may be created for legacy stock, but current mandates apply primarily to products placed on the market once the rules take effect, and carry no 3D geometry. **ReUseBIM** operates on this market stock. In addition, spreadsheet-based documentation schemes such as **Cirkla** [2] illustrate the burden of non-automated workflows, requiring up to 100 attributes per component.

Within visualization research, circular construction has received limited attention. Cosma et al. present an interactive tool for building-materials recovery [4], Preut et al. characterise circular-economy stakeholders’ information needs [15], and Hoffmann et al. design digital data streams for circularity in 3D printing [10]; urban material stock has been mapped for reuse at city scale [17]. Our work complements these by treating the reclaimed component as the unit of visual comparison and carrying it into the **BIM** authoring environment.

None of these approaches addresses both discoverability across multiple **DMBs** and **BIM**-ready integration of reclaimed components, which is the gap **ReUseBIM** aims to close.

3 SYSTEM

ReUseBIM has two parts, designed around the daily workflow of practising architects.

The **dashboard** (Figure 3) scrapes part listings - dimensions, cost, CO₂ equivalent, and attached files - from **DMBs** and presents them as a searchable, filterable card grid (**R1**). Each card exposes the metadata an architect needs to assess a component at a glance and provides a ready-made **IFC** file that can be dragged directly into any **BIM**-capable software (**R2**). Cards are ordered by recency of ingestion by default and can be re-sorted or filtered by category, dimensions, cost, or CO₂ equivalent. A background process auto-generates that **IFC** geometry from the stored dimensions, classifying each component into the appropriate entity type (e.g. **IfcBeam**, **IfcWall**, **IfcWindow**). A drag-and-drop workspace (cf. Figure 4) lets architects assemble a shortlist of candidates and download them in bulk. The **DMB** we used for demonstration purposes is [1], identified through the **Opalis** directory [14]; further sources can be added by implementing a thin scraper adapter per platform.

The **IFC** geometry is generated from the three stored dimensions (length, width, height) using an extruded area solid, conforming to the **IFC2X3** schema. Input data is whatever the **DMB** exposes: typically dimensions scraped from the product page, supplemented by any

attached CAD drawings or images. The system does not attempt to reconstruct detailed geometry from images or drawings; where only a bounding box can be derived, that is the geometry that is exported - a limitation discussed in Section 4. Entity type assignment is rule-based: components whose category contains “window” are mapped to IfcWindow, those containing “door” to IfcDoor, and all remaining structural components to IfcWall or IfcBeam depending on their aspect ratio. This is intentionally conservative - the goal is a geometrically correct bounding volume that places the component in the correct location within a BIM model, not a detailed parametric representation. Bounding-box correctness is verified automatically after generation by reading back the axis-aligned bounding box of the exported solid and comparing it against the stored dimensions.



Figure 3: A single card element of a reclaimed component with extracted metadata, cost, CO₂, and access to Building Information Modelling (BIM)-ready exports. It is readily available for drag-and-drop into the workspace and suitable Computer-Aided Design (CAD) software. Multiple media items can be uploaded and attached, also images. Photograph from [1].

The **enrichment pipeline** addresses the incompleteness that our DMB screening revealed (R3). The architect drops a target IFC file and a supporting document (PDF, spreadsheet, or image) onto the interface; the pipeline extracts text and prompts an LLM to identify attributes such as fire resistance class and thermal transmittance. A typical example is a CAD wireframe drawing of a reclaimed aluminium window supplied by the DMB alongside the listing. Detected values are shown to the architect for confirmation before being written into a dedicated IFC property set, leaving original geometry and existing property sets untouched. The extracted attributes are stored under a dedicated property set so that they remain distinguishable from manufacturer-supplied data and can be filtered or audited independently within the BIM model. The demonstrator parses the attached PDF in the browser and queries the mistral-medium model with a (German and English) prompt that fixes the output to one line per attribute [8], with an example given in the appendix (Prompt 1).




Figure 4: Drag-and-Drop Workspace for keeping track of live-use components, with bulk-downloading functionality. Photographs from [1].

We screened multiple DMBs, e.g. Concular, Baukarussell, re:store, Use Again, and Juch Areal [1], against decision-relevant attributes (e.g. weight, fire rating, load-bearing capacity) and found completeness varied sharply between platforms, motivating the enrichment pipeline as a complement to scraped data rather than a replacement for it [7]. Most platforms exposed basic geometric dimensions and a price, but technical attributes required for regulatory compliance - fire resistance class, thermal transmittance, load-bearing classification - were rarely present as structured fields and instead appeared, when available at all, only in attached PDF documents. Component condition and durability (e.g. corrosion, weld quality, or material aging for steel elements) were not assessed by any platform screened; this remains an open challenge for reuse decision-making that the current system does not address. The underlying data model is shown in Figure 5. Listings are stored relationally, with each part holding the fields taken directly from the source listing, namely name, dimensions, cost, and CO₂ equivalent, alongside the attached images and documents as retrieved. Everything else is derived by the system, since the IFC geometry is generated from the stored dimensions and the enriched attributes are extracted from those attachments, so provenance stays visible per field. The enrichment step is split across client and server, parsing the attached document to text in the browser and creating and relating the property set server-side through IfcOpenShell, following the architecture set out in our earlier work [7, 8].

4 DISCUSSION

The climate benefits of material reuse depend not only on the availability of reclaimed components but also on designers’ ability to efficiently identify, compare, and integrate them into ongoing projects. Our collaboration with practising architects suggests that, beyond material availability, the effort required to locate, evaluate, document, and integrate reclaimed components often discourages their use within the time constraints of conventional design practice. By reducing overhead through aggregated search (R1), metadata filtering (R1), one-click IFC export (R2), and drag-and-drop assembly (R2), the tool primarily targets the buildings domain and remains highly relevant to supply chains as well as cities and urban planning, particularly where municipal reuse mandates exist.

The current tool is a working proof-of-concept - a functional implementation demonstrating the proposed workflow, not yet a production-ready system. Each component is handled as an isolated IFC file, which does not scale to comparing many components within a single shared model; integration with a Common Data Environment such as Speckle [16], together with batch processing of many components




Figure 5: Data model of the [ReUseBIM](#) dashboard. The central Part model is outlined in teal. Each arrow denotes a one-to-many relationship.

at once, is needed for production use. Scraped listings currently pass through no automated quality control, so heterogeneous or malformed source data reaches the database unchecked and has to be corrected by hand. Attribute extraction (R3) relies on an online LLM and covers only a small set of attributes, raising data-privacy questions for proprietary documentation. Geometry reconstruction is limited to bounding-box extrusion from stored dimensions; components for which no reliable dimension data can be scraped cannot be represented in IFC. Component condition and durability - factors that are critical to the reuse decision for structural elements - are not assessable from the available metadata and would require physical inspection protocols beyond the scope of the current system. All of these are tractable next steps. In the interim, scraped entries can be corrected manually through the database editing interface shown in Figure 6.

An LLM-based pipeline aimed at reducing emissions consumes energy and water of its own [11]. Our usage stays deliberately narrow, with one short extraction call per document, set against a saving measured in the embodied carbon of a reused component. Whether efficiency gains of this kind hold at scale is contested [12], and quantifying both sides of this balance is left to future work.

The card grid supports search and filtering, but comparison across candidates remains manual. Visual comparison of metadata across shortlisted components, including location, condition, and delivery terms, would let architects weigh trade-offs that currently require inspecting each listing individually.

The work was developed iteratively together with domain experts from architecture practice with backgrounds in circular construction, whose feedback shaped the workflow design. Extraction accuracy has not been quantified, as that would require a labelled corpus of component documentation that does not yet exist. A formal evaluation with a broader panel of expert users - measuring task completion time, error rate, and perceived usefulness against the current manual workflow - remains to be conducted and is a prerequisite for stronger claims about usability and real-world impact.

5 CONCLUSION

Developed with practising architects, [ReUseBIM](#) addresses the two workflow obstacles that prevent building-material reuse from becoming routine today: scattered discoverability and incomplete metadata. The dashboard gives architects a single searchable view across DMBs, with drag-and-drop IFC exports ready for their BIM software. The enrichment pipeline closes data gaps by extracting missing attributes from unstructured documentation via an LLM, keeping the architect in control of every value written to the model.

The construction sector cannot meet net-zero targets if reuse remains a niche specialist activity. Making it as easy as specifying new materials - within the tools architects already use - is a concrete, near-term lever. [ReUseBIM](#) is a step in that direction, and we hope it prompts further work at the intersection of visual analytics, BIM workflows, and the circular built environment.

APPENDIX

Change Part

Aluminiumfenster 722x2082

Name:	Aluminiumfenster 722x2082
Slug:	aluminiumfenster-722x2082
Category:	Tragwerk ✎ + ✖ 👁
Description:	Aluminiumprofil (Legierung AL-MG-Si 0.5 F 22-28)
Length:	722.0
Width:	2082.0
Height:	91.0
Co2 eq:	
Cost:	
Site:	✎ + 👁
Original link:	Currently: https://juchareal.store/components/80-aluminiumfenster-722x2082 Change: https://juchareal.store/components/80-aluminiumfenster-722x2082
Meta data:	["Beschreibung": "Aluminiumprofil (Legierung AL-MG-Si 0.5 F 22-28)", "Dimensionen": "722x2082x91 mm", "Menge": "15 Stk", "Downloads": "DWG-21, IFC-21-1", "Uf-Wert-Rahmen z": "0.92 W/(m²K)", "Produzent": "Schüco", "Typ": "Schüco AWS/ADS 75-Si Standard eckig Typ A", "Schalldämmung RwP max": "48 dB(A)", "Glasaufbau": "3-fach-Glas", "UG-Wert Glas z": "0.5 W/(m²K)"]

Figure 6: Manual database editing interface to correct scraped entries. Similar views exist for all referenced database models (cf. [Figure 5](#)).

From the following building component spec sheet text, extract the values for 'Feuerwiderstand' or 'Feuerschutz' (fire resistance) and 'Wärmedämmung' (thermal insulation) in W/(m²K). Format exactly as:
Feuerwiderstand/Feuerschutz: [value]
Wärmedämmung: [value]
If a value is not found, use 'Not found'. Ignore all other information.

Prompt 1: Example prompt used by the enrichment pipeline [8]. Fire resistance and thermal transmittance were chosen because they determine whether a reclaimed component is admissible for reuse.

REFERENCES

- [1] “Recyclingzentrum Juch-Areal, Bauteilkatalog.” <https://juchareal.store/> (accessed June 30, 2026).
- [2] S. Ackermann and D. Vittani, “Cirkla Swiss Inventory Guidance v1.0,” 2026. <https://www.cirkla.ch/> (accessed June 26, 2026).
- [3] J. Brütting, J. Desruelle, G. Senatore, and C. Fivet, “Design of Truss Structures Through Reuse,” *Structures*, vol. 18, pp. 128–137, Apr. 2019, doi: [10.1016/j.istruc.2018.11.006](https://doi.org/10.1016/j.istruc.2018.11.006).
- [4] I. Cosma, I. Petre, C. Stalidi, G. Suci, and A. Dănilă, “An interactive visualization tool for circular economy and building materials recovery,” in *2024 IEEE 30th International Symposium for Design and Technology in Electronic Packaging (SIITME)*, Oct. 2024, pp. 251–254, doi: [10.1109/siitme63973.2024.10814893](https://doi.org/10.1109/siitme63973.2024.10814893).
- [5] European Commission, “Ecodesign for Sustainable Products Regulation,” June 2026. <https://eur-lex.europa.eu/eli/reg/2024/1781/oj/eng>.
- [6] L. Gschweidl, M. Lang-Raudaschl, D. Daleyev, and M. Stavric, “ReUse with BIM: Connecting architectural salvage centres to your BIM software,” in *Proceedings of the 42nd Conference on Education and Research in Computer Aided Architectural Design in Europe, eCAADe 2024*, Sept. 2024, vol. 2, pp. 9–16, doi: [10.52842/conf.ecaade.2024.2.009](https://doi.org/10.52842/conf.ecaade.2024.2.009).
- [7] L. Gschweidl, P. Waldert, M. Lang-Raudaschl, and M. Stavric, “Connecting Material Banks to BIM,” *Accepted Manuscript, submitted to the Net-Zero Future Conference*, 2026.
- [8] L. Gschweidl, P. Waldert, M. Lang-Raudaschl, and M. Stavric, “From Material Passports to Open BIM,” *Accepted Manuscript, submitted to eCAADe*, 2026.
- [9] F. Heisel and J. McGranahan, “Enabling Design for Circularity with Computational Tools,” in *A Circular Built Environment in the Digital Age*, Cham, Switzerland: Springer, 2024, pp. 97–110.
- [10] S. Hoffmann, N. W. Böcking, A. N. Shirazi, W. Schaefer, T. Richter, and V. Wulf, “Design and Development of Digital Data Streams for a Sustainable Circular Economy in 3D Printing,” in *Proceedings of the 12th International Conference on Communities & Technologies*, July 2025, pp. 356–370, doi: [10.1145/3742800.3742837](https://doi.org/10.1145/3742800.3742837).
- [11] N. Jegham, M. Abdelatti, C. Y. Koh, L. Elmoubarki, and A. Hendawi, “How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference,” 2025. <https://arxiv.org/abs/2505.09598>.
- [12] A. S. Luccioni, E. Strubell, and K. Crawford, “From Efficiency Gains to Rebound Effects: The Problem of Jevons’ Paradox in AI’s Polarized Environmental Debate,” in *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, June 2025, pp. 76–88, doi: [10.1145/3715275.3732007](https://doi.org/10.1145/3715275.3732007).
- [13] A. Marchal, “Basel’s Reuse Mandate and the Path to Net Zero 2037,” 2023.
- [14] Opalis, “Opalis: Material Bank Directory for France and Benelux,” 2026. <https://www.opalis.eu/> (accessed June 26, 2026).
- [15] A. Preut, T. Mertz, L. Cibulski, and J.-P. Kopka, “A User-Centered Perspective on Information Needs of Stakeholders in the Circular Economy,” in *2024 Electronics Goes Green 2024+ (EGG)*, June 2024, pp. 1–9, doi: [10.23919/egg62010.2024.10631223](https://doi.org/10.23919/egg62010.2024.10631223).
- [16] Speckle Systems, “Speckle: Open Data Platform for AEC,” 2026. <https://speckle.systems/> (accessed Aug. 11, 2026).
- [17] G. Suci, I. Petre, and C. Stalidi, “Urban transformation processes and material stock techniques to facilitate circular economy and urban resilience,” in *2023 15th International Conference on Electronics, Computers and Artificial Intelligence (ECAI)*, June 2023, pp. 1–5, doi: [10.1109/ecai58194.2023.10194029](https://doi.org/10.1109/ecai58194.2023.10194029).
- [18] B. Wibranek and O. Tessmann, “Augmented Reuse,” in *Towards Radical Regeneration*, Cham, Switzerland: Springer, 2022, pp. 411–423.
- [19] S. Çetin, C. De Wolf, and N. Bocken, “Circular Digital Built Environment: An Emerging Framework,” *Sustainability*, vol. 13, no. 11, p. 6348, June 2021, doi: [10.3390/su13116348](https://doi.org/10.3390/su13116348).