MULTIDIMENSIONAL WEIGHTED RANKING RECOMMENDATION SYSTEM FOR CIRCULAR MATERIAL MATCHMAKING

Noah Gerber[1], Ibrahim Al-Hazwani[2], and Pei-Yu Wu [3]
[1]ETH Zurich, Zurich, Switzerland
[2]University of Zurich, Zurich, Switzerland
[3]Delft University of Technology, Delft, The Netherlands

Abstract

Circular construction relies on digital marketplaces to overcome supply chain fragmentation, yet manual search for reclaimed components remains a critical bottleneck. This research develops a context-aware recommendation system for reusable material matchmaking. Utilizing a multidimensional weighted ranking algorithm with hard and soft constraints, the system adaptively aligns material attributes with project requirements. Validated through expert elicitation and a design simulation user study, results demonstrate high material findability and low cognitive load. However, insufficient algorithmic transparency can induce user friction. This work establishes a foundation for role-centric matchmaking, with future research focusing on similarity score cutoffs, enhanced interpretability, and API-driven cross-platform interoperability.

Introduction

Digital material marketplaces are evolving into mediatised platforms that foster a circular construction sector by enabling the sustainable reuse of building materials. Nevertheless, many existing reuse platforms operate within isolated local communities, preventing economies of scale and restricting the volume of available secondary products compared to primary materials. Consequently, the low quantity and high diversity of reclaimed batches often preclude their procurement in new construction projects. Recent literature identifies several critical barriers to platform effectiveness: (1) structural and economic obstacles such as high upfront costs; (2) operational constraints caused by the lack of spatial and temporal synchronization between deconstruction and construction phases; (3) informational gaps stemming from a lack of automated data entry and standardized performance documentation (Wilson et al., 2025; Quertamp et al., 2025); and (4) low user adoption. Notably, 70% of marketplaces prioritize auxiliary services (e.g., building inventory, logistics) over core digital infrastructure (Menny et al., 2024), resulting in underutilized databases that fail to reach the critical mass necessary to bridge the gap between supply and demand.

To address these challenges, research suggests transitioning from project-centered to role-centric matchmaking through asynchronous loops that mitigate strict synchronization requirements (Menny et al., 2024). Technological interventions such as blockchain, the InterPlanetary File System (IPFS) (Protocol Labs, 2026), and Digital Product Passports (Honic et al., 2024) are being explored to ensure traceability (Wilson et al., 2025). Furthermore, the integration of standardized circularity indicators and interoperable APIs aligned with Circular Economy standards (ISO/TC 323, 2018) facilitates cross-platform communication. Finally, participatory action research and regulatory alignment with Extended Producer Responsibility laws ensure that digital tools meet real-world operational needs and compliance requirements (Podkalicka and Fredriksson, 2023; Quertamp et al., 2025).

While these frameworks provide the necessary infrastructure for data exchange, their utility depends on early-stage intervention within the design process to maximize reuse potential before in-situ reclamation. Building upon this requirement, this study develops a Context-Aware Recommendation System (CARS) designed to facilitate the matchmaking of both in-situ components and off-site reclaimed materials. By employing a multi-phased participatory design approach with actors from Swiss marketplaces, this research seeks to: (1) identify critical decision-making factors to encode heterogeneous inputs into a unified schema; (2) integrate temporal and geospatial preferences into a context-aware engine to enhance logistical viability; and (3) provide explainable recommendations to foster user trust. To achieve this, the system encodes semantic and quantitative material records into a multidimensional vector database, facilitating similarity matching based on the mathematical proximity between dynamic user requirements and material attributes. By incorporating image data and explainability features, the system addresses the "black box" friction common in traditional vector retrieval.

The paper demonstrates the feasibility of integrating a CARS into a user-friendly interface that delivers personalized matchmaking based on dynamic constraints. By enabling users to interactively adjust priority levels across multilingual inventories, the system generates ranked outcomes with matching scores that reveal the underlying weighted logic. This guides users to iteratively refine search parameters and explore alternative material choices when exact matches are insufficient. Preliminary findings validate that high operational speed and intuitive weighted search logic effectively bridge the gap between technical retrieval and architectural design rationale. However, the results also highlight a critical need for enhanced semantic transparency and similarity thresholds to ensure reliability as databases scale. This research contributes a methodological framework for optimizing the quality and diversity of early-stage matchmaking in circular construction.

Related work

Relevance of context-aware recommender systems
CARS address the structural fragmentation and logistical constraints of material reuse by synchronizing spatial and temporal data. Unlike static systems, CARS incorporate dynamic attributes, such as disassembly windows and geographic proximity, to adapt results to immediate project needs (Adomavicius et al., 2011; Afzal et al., 2024). This approach mitigates the data sparsity and cold-start problems inherent in unique, transient secondary material databases by leveraging contextual constraints to estimate preferences when historical interaction data is absent (Kulkarni and Rodd, 2020). Consequently, the system models evolving professional requirements throughout the project lifecycle rather than relying on fixed interest models (Huang et al., 2024).

Distinctive nature of secondary matchmaking
Secondary material matchmaking differs from standard e-commerce due to high item heterogeneity and the absence of reciprocal supplier constraints. The process shifts from exact stock keeping unit matching to similarity matching of non-standardized components using quantitative, qualitative, and vector-based image data (Adomavicius et al., 2011). Furthermore, professional selection requires a multi-criteria evaluation framework rather than a single consumer rating (Afzal et al., 2024). Practitioners must concurrently weigh technical specifications, aesthetics, and environmental impact. This necessitates a ranking system capable of balancing diverse factors against specific professional workflows rather than simple popularity metrics.

Explainability in professional decision-support
Explainability transforms the recommendation engine from an opaque "black box" into a transparent support framework, establishing the epistemic trust necessary for high-stakes construction decisions (Tintarev and Masthoff, 2015). By allowing users to explicitly define attribute weighting, such as prioritizing structural load over aesthetics, the system provides requirement-specific justifications for its suggestions (Lai et al., 2024). This transparency is vital for validating unique, reclaimed components. Explicitly describing how a material aligns with complex project constraints ensures that recommendations are both actionable and interpretable within professional discovery processes.

Methodological framework

This research employs a multi-phased participatory design approach to develop a context-aware matchmaking platform. Initial scoping involved consultations with experts from Swiss secondary material companies to elicit operational workflows and define technical metadata requirements. An iterative development phase followed, where stakeholders evaluated prototype design considerations and prioritized functionalities. Finally, the platform underwent empirical evaluation through a structured user study with four native German-speaking participants. This linguistic selection ensured valid interaction with the localized interface and database.

User requirements

Technical and metadata requirements
Effective matchmaking necessitates the integration of high-fidelity material attributes across three categories: fundamental dimensions and classification; performance metrics (e.g., thermal or acoustic properties and mounting specifications); and a standardized physical condition scale. Expert elicitation indicates that practitioners prioritize chronological age over manufacturer identifiers when assessing reclaimed value. Furthermore, high-quality visual documentation is critical for the remote assessment of unique components. To ensure interoperability, the search schema aligns with DIN SPEC 91484 (DIN German Institute for Standardization, 2024) and Digital Product Passport frameworks (Honic et al., 2024), facilitating a unified infrastructure for cross-platform exchange.

Role-specific functional requirements
Matchmaking architectures must accommodate the divergent workflows of specialized professional groups. Architects and designers prioritize reliability and contextual inspiration, typically matching available inventories against pre-defined design constraints. Conversely, business owners and contractors emphasize carbon savings and regulatory compliance over initial procurement costs. To ensure functional integrity, these stakeholders require integrated services such as quality assurance, professional testing, and specialized cleaning. The system addresses these needs through a dual-announcement model comprising both supply listings and search requests. This bifacial architecture facilitates a proactive procurement process tailored to the logistical demands of circular economy actors.

Operational and logistical requirements
Systemic viability depends on managing the physical constraints of the construction sector. Geospatial filters are integrated to mitigate transport costs, which frequently dictate the economic feasibility of secondary materials. To address the primary cost drivers in the reuse cycle, the system supports material demand prediction to optimize storage management. Additionally, operational data regarding deconstruction viability (e.g., ease of disassembly, fragility, and weight) is essential for sourcing materials in situ.

Psychological and soft factors
Material provenance, historical context, and environmental significance are decisive factors in the selection of reclaimed components. To manage the complexity of

non-standardized items, the system utilizes intuitive interface features designed to minimize cognitive load and user anxiety. Successful matchmaking requires a synthesis of a component's history, physical health, and spatiotemporal availability. Aligning these narrative qualities with rigorous technical constraints ensures a reliable material exchange.

Design goals

Based on the research framework and expert elicitation, five core design goals are established to transition digital marketplaces toward multidimensional decision support.

1. Multidimensional similarity matching
The system implements similarity matching to shift from rigid queries to the flexible discovery of heterogeneous components. Representing qualitative and quantitative dimensions within a vector database enables the engine to suggest functional alternatives aligned with latent project objectives rather than restricted search terms.

2. Spatiotemporal and logistic optimization
Geospatial proximity and temporal availability are integrated as dynamic factors managed through user-defined priority levels. This addresses the critical synchronization required between in-situ deconstruction windows and reconstruction schedules, ensuring logistical feasibility.

3. Role-centric recommendation logic
The engine adapts weighting parameters to the distinct priorities of stakeholders. While individual users may focus on cost, professional architects often emphasize carbon savings and technical reliability. Tailoring the ranking logic to these requirements ensures professional relevance.

4. Trust through explainability and user experience design
Transparent recommendations detail the reasoning behind matches, such as shared mounting specifications or condition ratings. This transparency, combined with streamlined interaction patterns, bridges the trust gap regarding the performance of reclaimed materials in a risk-averse sector.

5. Data interoperability and aggregation
The platform functions as a meta-aggregator to address market fragmentation. By aligning with international standards and aggregating data across multiple engines, the system provides a comprehensive regional overview, ensuring recommendations are derived from high-quality, interoperable datasets.

System overview

Material data schema

The data integration and structural representation, denoted as M, synthesizes metadata of each material item into an integrated schema.

$$M = \{\mathrm{id}, I, Q, P, S, L, T\}$$

$$\begin{cases} I = (T, D) & \text{Title and description} \\ Q \in \{\text{new, good, poor}\} & \text{Quality interval} \\ P \in \mathbb{R}_{\geq 0} & \text{Unit price} \\ S = (w, h, d) \in (\mathbb{R}^{+})^3 & \text{Volumetric Dimensions} \\ L = (\phi, \lambda) & \text{Geospatial coordinates} \\ T = (t_{start}, t_{end}) & \text{Temporal period} \end{cases} \tag{1}$$

The material schema is formalized as a multi-layered tuple. The identification layer $\{\mathrm{id}, I\}$ integrates a unique identifier with a descriptive metadata. Economic and qualitative dimensions are represented by unit price P and material integrity Q, with the latter categorized into three discrete intervals to facilitate standardized circular recovery assessment. Finally, the system encodes physical, geospatial, and temporal constraints via S, L, and T, which define volumetric attributes (w, h, d), coordinates (ϕ, λ), and the operational validity period (t_{start}, t_{end}).

Platform architecture

The platform is engineered for data ingestion and semantic retrieval using a serverless framework. The system is partitioned into two distinct operational pipelines leveraging edge computing to minimize latency and ensure linear scalability, as illustrated in Figure 1.

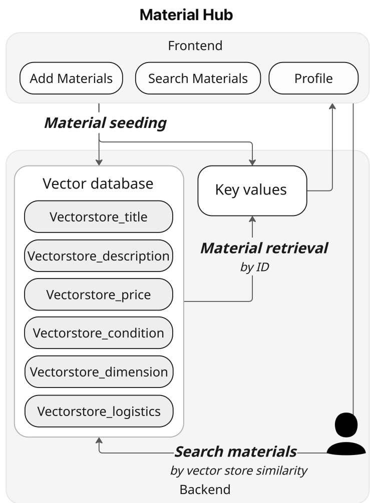

Figure 1: System archtiecture and workflow.

Material seeding pipeline

The seeding pipeline initiates with edge-based input processing via Cloudflare Workers. These functions execute data normalization and validation of attributes, including Swiss eBKP classification codes (Swiss Society of Engineers and Architects, 2020), pricing, and dimensions. Following refinement, the Gemini Embedding model (Lee et al., 2025) transforms textual descriptions and technical parameters into high-dimensional vectors. These representations are indexed within Vectorize, a distributed vector database maintaining the latent space of all material profiles.

Material retrieval pipeline

The retrieval pipeline facilitates semantic search by translating natural language queries into vector representations. The architecture executes similarity analysis using a k-nearest neighbor (k-NN) algorithm to calculate the cosine similarity between the query vector and stored indices. By identifying elements based on spatial proximity in the latent space, the system overcomes the limitations of traditional keyword matching for non-standardized material descriptions.

Material recommendation system widget

Integrating specifications into the digital *Material Hub* (Figure 2) triggers a ranked list of recommendations based on calculated matching scores. System widgets allow for the iterative adjustment of preference weights, enabling real-time exploration of alternative candidates.

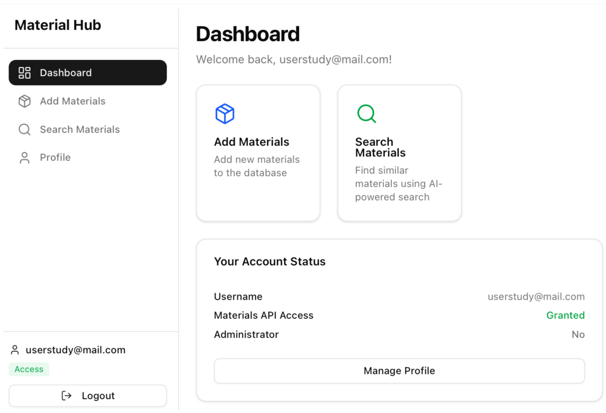

Figure 2: Material Hub dashboard prototype highlighting the integrated modules for material registration, multidimensional search, and user profile management.

Material data ingestion

The system supports dual-mode data ingestion: automated batch processing via inventory files (e.g., pre-renovation or pre-demolition audits) or manual entry, as shown in Figure 3. The automated pipeline utilizes Large Language Models (LLMs), specifically Llama 3.1 8B Instruct (Meta AI, 2024), to parse unstructured data and extract material properties. In cases of data scarcity, the system maintains integrity by imputing missing values based on category averages. Imported materials are automatically mapped to Swiss eBKP-H classification codes (Swiss Society of Engineers and Architects, 2020) to standardize heterogeneous documentation into a structured database.

Preference adjustments

To optimize discovery within multidimensional decision spaces, the system utilizes a weighted vector retrieval

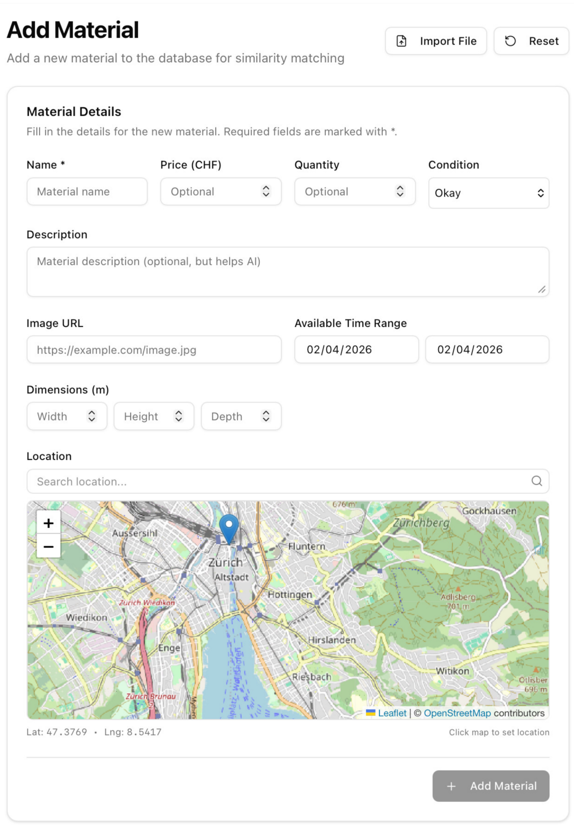

Figure 3: User interface for material data ingestion by batch processing or manual entry.

mechanism. Users calibrate the recommendation engine by assigning discrete priority levels to specific attributes, which are translated into quantitative weights. This process employs a dual-constraint logic: hard constraints serve as binary filters to exclude non-compliant materials, while default soft constraints facilitate fuzzy matching to preserve high-potential alternatives.

Geospatial logistics and proximity

The system integrates an interactive mapping interface powered by OpenStreetMap to manage geographical constraints. Users define a central project coordinate and a search radius, restricted to a maximum of 100 km. This parameter is empirically derived from expert interviews, reflecting the logistical threshold where transportation costs and environmental impacts typically become counterproductive for reclaimed materials.

Global weight distribution and ranked results

To enhance interpretability, the system visualizes the contribution of individual attributes to the final similarity scores. As illustrated in Figure 5, the global weighted attributes of the aggregated k-matched items are displayed in a stacked bar chart.

The top k items are ranked using a late-fusion integration logic that synthesizes unstructured semantic data with structured constraints, presented in modular cards in Figure 6. To provide granular transparency, each entry dis-

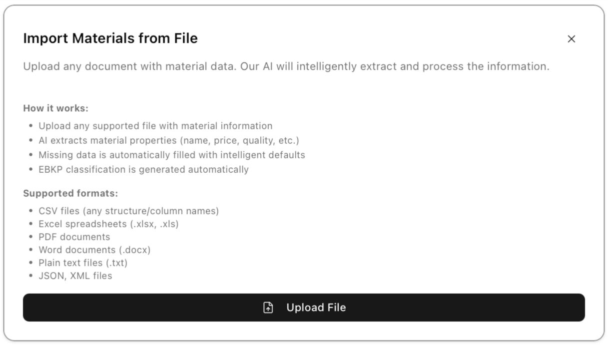

Figure 4: Material information import, extraction, and standardization.

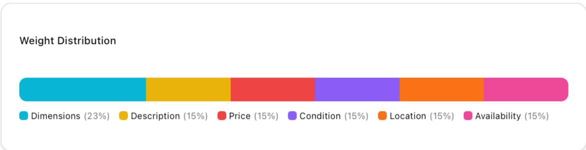

Figure 5: Global weight of attribute-level contribution to the similarity scores of k-matched items.

plays a "local weight" representing the aggregate alignment of its specific attributes. We define $s_{\text{sem}}(m_i)$ as the cosine similarity between embeddings, and $s_j(m_i) \in [0, 1]$ as normalized scores for J structured preference criteria (e.g., distance, quantity, and condition). The final ranking score $S_{\text{final}}(m_i)$ is produced via linear aggregation: Items are sorted in descending order of $S_{\text{final}}(m_i)$. This framework ensures the ranking remains interpretable, allowing users to discern the relative importance of latent semantic relevance versus explicit attribute requirements.

$$S_{\text{final}}(m_i) = w_{\text{sem}}\, s_{\text{sem}}(m_i) + \sum_{j=1}^{J} w_j\, s_j(m_i),$$

$$w_{\text{sem}} + \sum_{j=1}^{J} w_j = 1, \quad w_{\text{sem}}, w_j \geq 0.$$

System evaluation

This section presents a multi-faceted evaluation of the platform, assessing technical efficiency via quantitative metrics, retrieval accuracy through a hybrid ranking analysis, and practical utility via controlled user studies.

Performance assessment

Performance was evaluated based on latency (end-to-end processing duration) and scalability (performance stability relative to repository growth). Latency measurements across the retrieval pipeline, comprising embedding generation, vector similarity search, deterministic filtering, and final scoring, indicate that response times remain within thresholds for real-time interaction. Vector similarity search is the most intensive stage. However, the

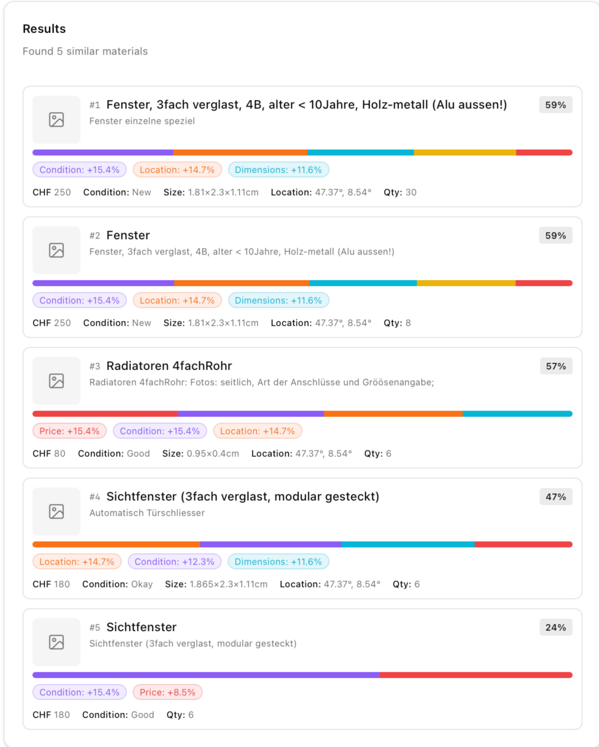

Figure 6: Ranked material candidates with localized attribute decomposition for the search query "windows."

implementation of approximate nearest neighbor indexing significantly reduces retrieval time compared to linear scans. Embedding generation and final ranking operations contribute marginal overhead, as the latter processes a pruned candidate subset.

Incremental dataset scaling experiments confirm that the approximate nearest neighbor indexing structure maintains sub-linear growth, avoiding exhaustive comparisons. This scalability and throughput is augmented by a serverless architecture that enables concurrent operations through automatic worker instantiation. These findings suggest the architecture can support expanding marketplaces without significant performance degradation.

Retrieval quality and error analysis

Preliminary quality assessments verified the alignment of retrieved materials with semantic intent and structured constraints. A comparative analysis between purely semantic and hybrid strategies reveals that while semantic-only retrieval excels at discovering conceptually related materials with divergent terminology, it frequently identifies candidates that fail physical or logistical requirements (e.g., geographic proximity). The hybrid pipeline resolves this by integrating deterministic filtering, successfully pruning unsuitable results while preserving the discovery advantages of vector search.

Error analysis identified three primary failure modes: (1) data sparsity, where incomplete descriptions prevent high-fidelity semantic representations; (2) category overlap, where distinct materials (e.g., panels vs. beams) share similar descriptive language; and (3) linguistic ambiguity, where vague queries produce broad, low-precision result sets. These findings underscore that structured metadata is essential for resolving the inherent ambiguities of natural language in technical domains.

User study

The user study evaluated the platform's functional performance across four dimensions: usability, retrieval efficacy, design alignment, and qualitative feedback. Participants conducted a design simulation for a temporary pavilion in Zurich, navigating a database to source technical components (e.g., glazing units, HVAC systems) under specific lifespan and safety constraints. This iterative process, as shown in Figure 7, facilitated a multidimensional analysis of system usability, the efficacy of the recommendation engine in discovery, and the alignment of results with architectural design rationales.

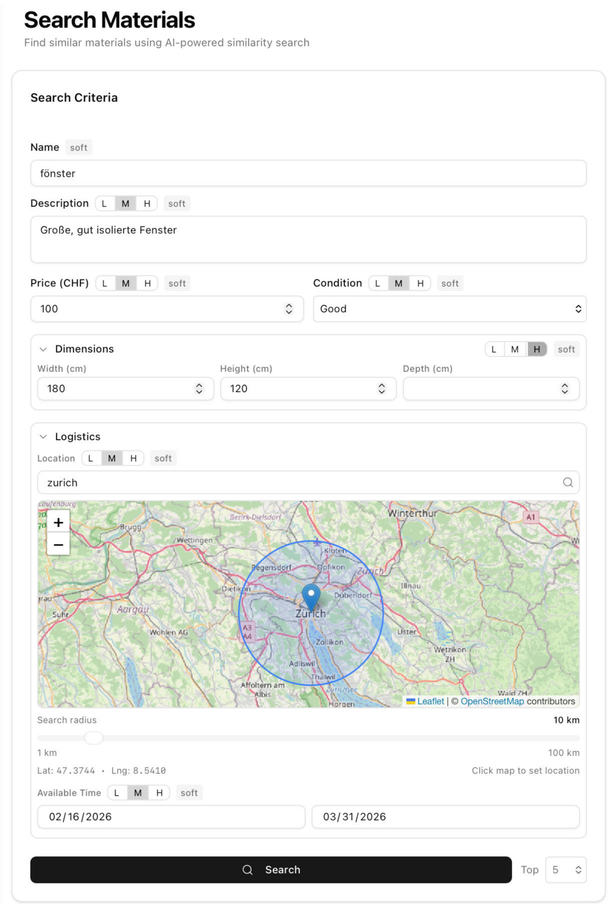

Figure 7: User interface for multidimensional vector search with constraint-based material retrieval logic.

System usability and user experience

Usability was assessed via metrics adapted from the System Usability Scale (Lewis, 2018). Participants reported low system complexity and high consistency, evidenced by high learnability scores (Median = 4.5) and low cognitive load. Technical confidence remained high (Median = 4.0), though variability (SD = 1.50) indicates that individual technical self-efficacy influences the user experience. Workflow fluidity analysis showed strong functional integration between vector search and filtering (Mean = 3.75, SD = 0.50). While 75% of participants perceived the system as highly consistent, divergent scores in intended frequency of use suggest that perceived utility is sensitive to individual professional context.

Information retrieval and discovery efficacy

The efficacy of the vector store and recommendation engine was validated by a mean search success score of 8.0 (SD = 1.15) on a 10-point scale, with no reported search failures. Participants identified ”description”, ”name”, and ”size” as the most critical metadata for procurement decision-making. However, a ”black-box” effect was observed regarding the ranking logic. Participants noted that weighted search results lacked sufficient descriptive interpretation, highlighting a requirement for greater algorithmic transparency and interpretability in the retrieval output.

Design rationale and requirement alignment

The simulation demonstrated the platform's ability to bridge technical retrieval with architectural application. Material selection was driven by semantic matches (type, color) and structured constraints (price, location, and age). Participants successfully identified components aligned with project requirements, demonstrating a correlation between search precision and the ability to justify circularity. However, search friction occurred when qualitative reasoning was required, suggesting that effective reuse advocacy depends on both engine precision and the user's prior technical knowledge.

Qualitative feedback and iterative improvement

Participants cited high operational speed and the weighting algorithm as primary drivers of efficacy. Interactive filters provided ”positive friction”, allowing users to refine discovery based on geographical requirements. Conversely, significant bottlenecks included high initial information density and the absence of ”empty state” feedback. The vector engine's tendency to return loosely related items for unmatched queries, such as displaying glazing units for sanitary ware, caused user frustration.

Discussions

Theoretical positioning and contributions

A primary challenge in circular construction platforms is the semantic gap caused by non-standardized, short-text descriptions, which often triggers cold-start failures in traditional algorithms (Van Capelleveen et al., 2021). While existing systems rely on rigid feature crosses and static metadata (Tserng et al., 2025), they struggle with the linguistic ambiguity of ”in-the-wild” data. Our research addresses this by implementing a late-fusion linear aggregation logic. By utilizing semantic embeddings to navigate jargon-heavy descriptions, the system identifies conceptually relevant materials that keyword-based searches exclude. This adaptive approach enables soft trade-offs where high semantic relevance compensates for minor attribute mismatches, ensuring findability despite the data sparsity identified in our error analysis.

Our evaluation confirms that algorithmic "black-box" behavior remains a significant source of user friction (Blackburn et al., 2026). Participants noted that while the ranking was successful, the rationale behind specific scores remained opaque. This validates the necessity of our proposed "glass-box" framework. By exposing ranking logic via local attribute decomposition, the system allows professional users to verify how semantic intent is balanced against physical requirements. This transparency is critical for reducing the cognitive load and decision-making friction observed during the user study, where a lack of descriptive interpretation was highlighted as a primary bottleneck.

The architecture is positioned to facilitate high-value component flows between demolition and construction, addressing the need for platforms to move beyond disciplinary silos (Blackburn et al., 2026). By adopting an API-ready framework, this work supports the platform-enabled coordination necessary to close resource loops. The multidimensional weighted ranking provides the governance logic required to manage ecosystem-wide interdependencies, shifting the focus from low-grade recycling to high-value resource retention and supporting circular construction through role-centric matchmaking.

Limitations and future work

The system evaluation revealed several technical and user experience limitations that define our future research pathway. The error analysis highlighted a critical need for a similarity score cutoff. Currently, the vector engine lacks a hard threshold, occasionally retrieving unrelated items (e.g., displaying glazing units for sanitary ware queries) when direct matches are absent. Implementing a calibrated cutoff will enable "no results found" notifications, preventing the user frustration observed during redundant filter refinement. While the weighting algorithm is technically sound, the user study identified a gap in semantic transparency. Future iterations will replace the current toggles with dynamic weight tuning and introduce semantic proximity labels. These features, alongside integrated imagery and empty-state feedback, will transform the ranking logic into a more interpretable tool, mitigating the information density issues reported by participants. To enhance data robustness, the architecture will be restructured as an API service layer to ingest structured material data from existing public marketplaces. Furthermore, the limited sample size necessitates future evaluations with larger, diverse participant groups to benchmark the hybrid system against other context-aware recommender systems and validate the findings across broader supply-and-demand scenarios.

Conclusions

This study validates that a multidimensional vector ranking system, augmented by dual-constraint logic, effectively supports the sourcing of reusable materials for architectural design. The preliminary findings indicate that while high operational speeds and weighted results facilitate discovery, the lack of transparency in algorithmic ranking and the ambiguity of weight-based toggles can obscure the system's underlying logic. Participants highlighted a "black box" friction where the rationale for specific material suggestions remained unclear, confirming that explainability is as critical as retrieval accuracy for professional adoption. To address these challenges, the research establishes that implementing similarity thresholds, clearer semantic labeling, and the progressive disclosure of filters is necessary to mitigate user frustration and build trust in the recommendation output. Future development will focus on API-driven scalability and the integration of broader marketplace datasets to solidify the platform's role as a meta-aggregator. By bridging the gap between technical material retrieval and architectural design rationale, this work provides a methodological foundation for optimizing circular material flows within the construction sector.

References

Adomavicius, G., Mobasher, B., Ricci, F., and Tuzhilin, A. (2011). Context-aware recommender systems. AI Magazine, 32(3):67–80.

Afzal, I., Yilmazel, B., and Kaleli, C. (2024). An approach for multi-context-aware multi-criteria recommender systems based on deep learning. IEEE Access, 12:99936–99948.

Blackburn, O., Ritala, P., Keränen, J., and Bocken, N. (2026). Circular economy platforms: A systematic review. Business Strategy and the Environment, 35(2):2810–2841.

DIN German Institute for Standardization (2024). Din spec 91481:2024-02 classification of recycled plastics of polyamides by data quality levels for use and (digital) trading.

Honic, M., Magalhães, P. M., and Bosch, P. V. D. (2024). From data templates to material passports and digital product passports. In A Circular Built Environment in the Digital Age, pages 79–94. Springer.

Huang, S., Yang, H., Yao, Y., Lin, X., and Tu, Y. (2024). Deep adaptive interest network: Personalized recommendation with context-aware learning. arXiv preprint arXiv:2401.00000.

ISO/TC 323 (2018). Circular economy: Standards and guidance.

Kulkarni, S. and Rodd, S. F. (2020). Context aware recommendation systems: A review of the state of the art techniques. Computer Science Review, 37:100255.

Lai, K.-H., Yang, Z.-R., Lai, P.-Y., Wang, C.-D., Guizani, M., and Chen, M. (2024). Knowledge-aware explainable reciprocal recommendation. In The Thirty-Eighth AAAI Conference on Artificial Intelligence.

Lee, J., Chen, F., Dua, S., Cer, D., Shanbhogue, M., Naim, I., Ábrego, G. H., Li, Z., Chen, K., Vera, H. S., Ren, X., Zhang, S., Salz, D., Boratko, M., Han, J., Chen, B., Huang, S., Rao, V., Suganthan, P., Han, F., Doumanoglou, A., Gupta, N., Moiseev, F., Yip, C., Jain, A., Baumgartner, S., Shahi, S., Gomez, F. P., Mariserla, S., Choi, M., Shah, P., Goenka, S., Chen, K., Xia, Y., Chen, K., Duddu, S. M. K., Chen, Y., Walker, T., Zhou, W., Ghiya, R., Gleicher, Z., Gill, K., Dong, Z., Seyedhosseini, M., Sung, Y., Hoffmann, R., and Duerig, T. (2025). Gemini embedding: Generalizable embeddings from gemini. Technical report, Google DeepMind.

Lewis, J. R. (2018). The system usability scale: Past, present, and future. International Journal of Human-Computer Interaction, 34(7):577–590.

Menny, T., Le Guirriec, S., and De Wolf, C. (2024). The butterfly matchmaking model for circular construction: Towards a digital matchmaking platform tailored to french policy. Sustainable Production and Consumption, 49:130–143.

Meta AI (2024). Llama 3.1 8B Instruct.

Podkalicka, A. and Fredriksson, M. (2023). Mediatised marketplaces: Platforms, places, and strategies for trading material goods in digital economies. Convergence, 29(5):1352–1368.

Protocol Labs (2026). IPFS - Powers the Distributed Web.

Quertamp, F., Tang, T. H. L., and Ly, V. K. (2025). Designing a digital marketplace for vietnam's plastic recycling sector: Needs assessment and strategic blueprint. Technical report, United Nations Development Programme in Viet Nam, Ha Noi.

Swiss Society of Engineers and Architects (2020). ebkp-h sn 506 511 elementbasierter baukostenplan hochbau.

Tintarev, N. and Masthoff, J. (2015). Explaining recommendations: Design and evaluation. In Ricci, F., Rokach, L., and Shapira, B., editors, Recommender Systems Handbook, pages 353–382. Springer, 2 edition.

Tserng, H.-P., Chang, C.-C., and Tai, S.-H. (2025). The development of building materials recommendation system based on feature crosses. Advances in Civil Engineering, 1.

Van Capelleveen, G., Amrit, C., Zijm, H., Yazan, D. M., and Abdi, A. (2021). Toward building recommender systems for the circular economy: Exploring the perils of the european waste catalogue. Journal of Environmental Management, 277(111430).

Wilson, S., Adu-Duodu, K., Li, Y., Sham, R., Wang, Y., Solaiman, E., Perera, C., Ranjan, R., and Rana, O. (2025). Trustworthy second-hand marketplace for built environment. Technical report, Technical Report 111.