architecture

MDPI

Article

Semantic Ambiguity in Japanese Architectural Terms: A Computational Contextual Analysis

Gledis Gjata [1,*] and Satoshi Yamada [2]

[1] Graduate School of Science and Engineering, Ritsumeikan University, Kusatsu 525-8577, Japan
[2] Department of Architecture and Urban Design, College of Science and Engineering, Ritsumeikan University, Kusatsu 525-8577, Japan; sy@fc.ritsumei.ac.jp
* Correspondence: gr0625rr@ed.ritsumei.ac.jp; Tel.: +81-80-9566-9189

Abstract

Japanese architectural terms are often translated into English as if they named stable spatial objects, yet many of these terms carry material, relational, symbolic, and experiential meanings. This study investigates contextual variation relevant to possible semantic ambiguity in English-language discussions of eight Japanese architectural terms: engawa, tokonoma, ma, mu, shakkei, doma, shōji, and fusuma. It uses an audited contextual-semantic workflow based on a fine-tuned multilingual BERT masked-language model and two paired corpus conditions. FULL preserves broader explanatory and glossary-like contexts, while CLEAN removes explicitly definitional sentences to test whether contextual patterns persist after definitional support is reduced. The analysis combines target screening, sentence-level contextual clustering, architectural register alignment, cluster-aware interpretation, and a focused concept-alignment test for engawa. The results show uneven evidence of recurring contextual differentiation across the target set; the clusters do not independently establish distinct lexical senses. Engawa, tokonoma, ma, shakkei, and shōji show the most stable contextual structure, although the semantic significance of their clusters varies; mu functions mainly as a conceptual comparison case, fusuma supports a cautious partition-and-recomposition reading, and doma remains a low-stability diagnostic case. For engawa, positive similarity to preselected boundary-edge and connector-flow prompts is compatible with a combined architectural reading, developed here as the modal littoral-circuit space formulation. Overall, the study shows that contextual modeling can make ambiguity in Japanese architectural language more explicit and testable without reducing these terms to fixed semantic categories.

Keywords: Japanese architecture; architectural terminology; architectural semantics; architectural discourse; architectural interpretations; semantic ambiguity; polysemy; contextual language model

Check for updates

Academic Editor: Avi Friedman

Received: 2 July 2026
Revised: 28 September 2026
Accepted: 1 October 2026
Published: 6 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Architectural language is not a neutral inventory of built objects. Architectural terms do more than denote physical elements; they also shape how we conceive, interpret, and value space. This broad view is well established in architectural thought. Forty [1] showed that architectural vocabulary is historically active, and Pérez-Gómez [2], Harries [3], and Pallasmaa [4] each, by different routes, treat architecture as more than physical form alone. Hillier and Hanson [5] provide a related spatial argument: their account treats the ordering of architectural space not merely as physical composition, but as a means

through which relations between people and forms of social organization are structured. Architectural terms can therefore appear materially concrete while also carrying conceptual, relational, experiential, or social force.

These issues are particularly relevant to the Japanese architectural discourse examined here, which provides the bounded case for studying how apparently stable spatial terms acquire recurring material, symbolic, relational, and experiential meanings. This scope does not imply that semantic ambiguity is unique to Japanese terminology. Amit [6] argues that contemporary Japanese aesthetics contains internal distinctions that are not always visible from terminology alone. Likewise, writing on Japanese architecture, space, and aesthetics has repeatedly emphasized enclosure, ritual, domestic ordering, and relations between spatial form and social practice instead of isolated objecthood alone [7–9]. Pilgrim's account of ma adds a more specific interval-based frame, treating ma as a spatial, temporal, relational, and religio-aesthetic condition rather than as a fixed object [10]. Similarly, Nitschke [11] links ma to place, interval, lived space–time, and relational spatial experience. Isozaki [12] sharpens a related point by treating Japanese architecture through temporality, ritual reconstruction, ma/interstice, textual interpretation, and architectural ambiguity. Work on Japanese aesthetics further suggests that aesthetic meaning is constructed through language, perception, representation, external reality, nature, and wider structures of significance, rather than through stable object reference alone [13]. Saito [14] extends this discussion toward ordinary life and ethical orientation. Taken together, these studies suggest that ambiguity in Japanese architectural language should not be treated only as unclear wording or imprecise translation. It may be one way through which certain terms gain architectural and theoretical significance.

The issue is especially visible in several of the focal terms examined in this paper. Shakkei has been treated as a compositional relation through which external scenery is incorporated into the designed field, and not as a simple decorative device [15]. Kuitert [16] extends this discussion by arguing that jiejing/shakkei should not be reduced to a single visual borrowing technique, but should be understood as a broader landscape-design relation in which the landscape itself "lends" to garden-making.

Engawa has likewise been read as more than a peripheral strip. Kurokawa's account of engawa as a third type of space gives this problem a sharper architectural anchor, since it frames engawa as neither simply interior nor exterior, but as an intermediary condition between house and garden [17]. Studies of engawa situate the term within circulation, spatial mediation, and relations between interior and exterior. Wong [18] frames this through ring-based circulation and inside-outside transition in Ando's museums, while Hou et al. [19] emphasize boundary permeability and sequential spatial mediation. Yang's study of shikii is relevant as a related threshold framework because it examines the negotiation of public, common, and private space, but it should not be treated as direct evidence about engawa [20]. Tokonoma complicates the issue in different ways, as it names a physical alcove while also carrying ritual, visual, and symbolic associations. Studies of tea ritual and tea space show that ritual action, display, calligraphy, flower arrangement, and symbolic placement can charge spatial elements with meanings that exceed simple physical location [21–23]. Stoichita's account of framed and embedded pictorial space provides a useful art-historical parallel, as framed surfaces can function not only as limits or enclosures but also as spaces that stage another representational field within the act of viewing [24]. Studies of Japanese domestic interiors further complicate fixed models of reception space. Ingarden and Watanabe [25] emphasize overlapping domestic and social codes, transitional elements, flexible partitions, and negotiated boundaries, while still treating hierarchy and formal reception roles as part of the domestic order. Kim [26], more specifically, complicates reception-centered hierarchy by showing that Edo middle- and lower-ranking samurai residences combine the Omote/Ura distinction with spatial continuity, intermediary zones, outward-oriented connections, and

integrated domestic-life centers. Nuijsink [27] supports the separate weak-enclosure point in contemporary Japanese housing, where domestic interiors are reworked through thin, perforated, loosely defined, or non-room-like divisions that permit separation and connection at once. Across the literature discussed above, several focal terms, especially shakkei, engawa, and tokonoma, combine material, spatial, symbolic, and relational meanings. These claims have usually been developed through qualitative interpretation, with limited reproducible corpus-based testing.

The objective of this paper is to examine whether recurring contextual variation in English-language uses of Japanese architectural terms provides evidence relevant to semantic ambiguity. It asks whether a small, screened set of terms shows recurring contextual patterns across two corpus conditions. More specifically, the paper uses a domain-adapted contextual semantic model to examine ambiguity after reducing overt definitional support. The analysis examines contextual variation in a mixed English-language corpus of explanatory sources and peer-reviewed architectural literature, without treating those source types as separate analytical conditions.

That objective addresses two linked gaps. The first gap is architectural: although scholarship on several focal terms in this study, especially shakkei, engawa, and tokonoma, is substantial, their ambiguity has rarely been examined as a distinct analytical problem. The second gap is methodological: computational corpus methods have rarely been used to audit semantically dense Japanese architectural terms. Computational text analysis has expanded in construction-related research, especially in NLP and text-mining work on construction documentation, safety management, compliance checking, BIM, and risk management [28,29]. Corpus-based and discourse-oriented approaches have also begun to appear in architectural research, particularly in studies of computational architecture and architectural discourse [30,31]. Broader digital-humanities work further shows the relevance of deep-learning-based text analysis for humanities and information-science text resources, while also emphasizing training-data and domain-adaptation constraints [32]. Yet these methods have been applied only rarely to semantically dense Japanese architectural terms in explicit corpus audits. This paper contributes by connecting architectural interpretation with a corpus-based method for testing recurring contextual ambiguity. The Japanese terminology examined here defines the study's empirical scope rather than the limits of the framework. With appropriate corpus construction, model adaptation, and architecture-facing interpretive criteria, the same analytical logic could be adapted to other terminological fields and linguistic or architectural contexts, although such applications would require separate validation.

A contextual semantic approach fits this problem because the study asks how representations and uses of a term vary across local contexts, not only how often the term appears. Transformer models introduced a self-attention architecture that represents tokens through relations across a sequence rather than through recurrent or convolutional structure [33]. BERT adapts this architecture as a bidirectional Transformer encoder that jointly conditions token representations on left and right context [34]. By contrast, static models, like Word2Vec, assign a single vector to a lexical type [35]. Research on vector semantics treats polysemy as a core problem for distributional meaning models because a single representation learned across contexts can conflate word senses [36].

In contrast, word-sense disambiguation research defines the task as identifying a word's meaning in context [37]. Later work on contextualized representations shows that models such as ELMo, BERT, and GPT-2 produce token representations that vary with context [38], while BERTology research further examines how BERT encodes linguistic, semantic, and world knowledge and how its contextual representations behave across lay-

ers [39]. For terms that name physical elements while also carrying symbolic or relational meanings, contextual models can identify multiple recurring uses.

Model choice alone leaves the methodological problem unresolved. In machine learning research, reproducibility is necessary to verify the reliability of reported findings and reduce avoidable workflow errors [40]. Here, reproducibility controls, corpus auditing, and preserved run records form part of the evidentiary basis for architectural interpretation. The analysis builds on three prior studies by the same authors. The first validates the audited framework used here and establishes the FULL versus CLEAN contrast as a definitional audit [41]. The second study expands that research program by quantitatively mapping conceptual hierarchies and data-driven taxonomies across 28 Japanese architectural concepts [42]. A third study extends the same research program through computational contextual comparison of selected architectural terms [43]. As in these previous studies, this paper uses the FULL and CLEAN corpora. FULL preserves the broader expository context in which the terms circulate. CLEAN removes glossary-like and explicitly definitional scaffolding to examine whether contextual patterns remain identifiable under tighter filtering.

The cited literature provides the interpretive frame for reading the computational results. The five registers used later in the analysis, physical–spatial, symbolic–theoretical, represented/imagined, boundary/transition, and hybrid, serve as limited analytical categories rather than as a complete semantic system.

Accordingly, the study answers four research questions:

- RQ1. Do selected Japanese architectural terms exhibit recurring contextual-semantic structures that support interpreting them as more than singular material labels?
- RQ2. Do these contextual-semantic structures remain recoverable across FULL and CLEAN corpus conditions after overt definitional support is reduced?
- RQ3. Can the identified patterns be interpreted through the architectural registers of physical–spatial, symbolic–theoretical, represented/imagined, boundary/transition, and hybrid without treating those registers as a complete semantic ontology?
- RQ4. In the case of engawa, does the focused concept-alignment test support an architectural interpretation that combines connector-flow logic with boundary-edge logic?

The remainder of the paper is organized as follows. Section 2 presents the corpus design, audit logic, and analytical procedure. Section 3 reports the ambiguity analysis results. Section 4 interprets the results, and Section 5 summarizes the findings, returns to the research questions, and clarifies what an ambiguity-focused architectural semantic analysis can add to architectural studies and architectural discourse.

2. Methods

The complete codebase, active YAML configuration files, logs, manifests, CSV outputs, and exported figures for the reported run are available in the public repository listed in the Data and code availability statement. The Section 2 reports the analytical design and run settings, while the repository provides the full implementation record.

2.1. Research Design

This study examines ambiguity in eight Japanese architectural terms by applying an audited contextual-semantic framework to two paired corpus conditions, FULL and CLEAN [41]. It uses that framework for a narrower architectural problem: whether local ambiguity within selected terms can be made more explicit, tested for recurring internal structure, and interpreted through the study's predefined architectural registers. The study is deliberately narrow, examining a screened set of architectural terms within a wider research

program on Japanese architectural vocabulary. Its outputs serve as evidence for architectural interpretation, not proof of final semantic categories.

The study is a narrowed ambiguity analysis within a wider research program [41–43]. Earlier work established and validated the computational framework, mapped a heterogeneous concept field spanning abstract notions, materially anchored architectural elements, and liminal or boundary-oriented cases, and extended the framework through contextual comparison of selected architectural terms. The present study narrows the analytical target to a small, eight-term, screened pool suited to local ambiguity analysis under paired FULL and CLEAN conditions.

The workflow comprised five linked stages: a screening audit, contextual cluster identification, architectural register alignment, cluster-aware interpretation, and a focused secondary concept-alignment module for engawa. The order of these five stages is methodologically important. The analysis first assessed whether each term had enough usable evidence. It then identified contextual clusters and aligned those clusters with architectural registers. The later cluster-aware stage used the identified cluster centroids and sizes to refine the interpretation, avoiding reliance on term-level averages alone. The architectural registers and their prompts were specified from prior architectural interpretation. Although the pipeline applied them after screening and clustering, their content guided the subsequent alignment and did not emerge from the clusters themselves. The fine-tuned model and paired FULL/CLEAN corpora are inherited unchanged from the preceding validated framework [41], while the five-stage workflow shown here constitutes the ambiguity analysis performed in the present study. The overall analytical sequence is summarized in (Figure 1).

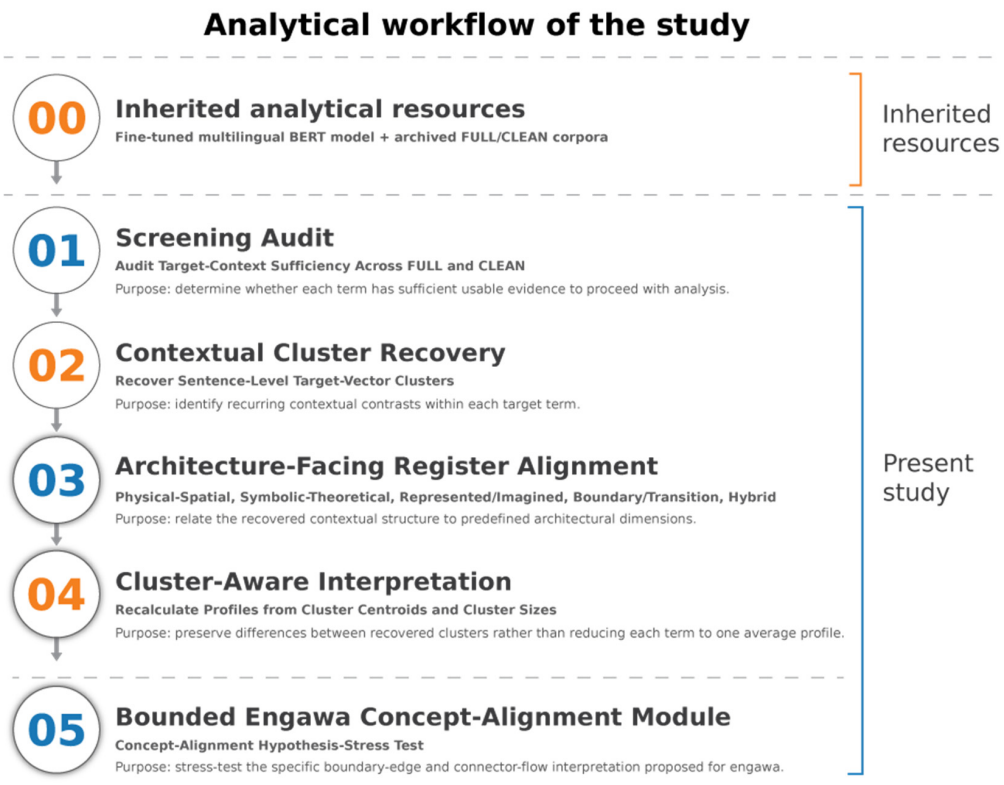


Figure 1. Analytical workflow of the study. The figure distinguishes between the inherited analytical resources used in the present study, namely the fine-tuned multilingual BERT model and the archived FULL/CLEAN corpora, and the five-stage ambiguity analysis performed here. Stages 01–05 summarize the analytical workflow and state each stage's purpose. The final engawa module is an exploratory, theory-guided concept-alignment probe applied only to engawa.

2.2. Corpus Basis and Paired Discourse Conditions

The present analysis draws on the corpus framework created in Gjata and Yamada [41] and reuses the same two audited discourse conditions, FULL and CLEAN. For the contextual BERT pathway used in this study, the FULL corpus contains 1,975,868 words and the CLEAN corpus 1,968,370 words. CLEAN was generated from FULL by removing glossary-style and explicitly definitional sentences, so the FULL-CLEAN comparison examines whether the target terms' contextual patterns persist after overt definitional support is removed from the same underlying source pool. The raw corpus, filtered corpora, source lists, and associated archived materials are publicly available through the repository released with Gjata and Yamada [41]. Briefly, the inherited corpus is an English-language expository corpus assembled primarily from web-based encyclopedic and related explanatory sources, supplemented by peer-reviewed architectural literature, as documented in the original corpus study [41]. The present study does not reconstruct or resample that corpus; it reuses the archived FULL and CLEAN versions so that the ambiguity analysis remains directly comparable with the preceding studies.

The inherited corpus includes broad expository writing on architecture, history, philosophy, aesthetics, and theology, as well as religious and ritual contexts that shaped parts of the wider corpus used in the earlier studies. This breadth allows the study to examine the terms across general explanatory material and the architectural literature included in the corpus, while leaving patterns specific to specialist writing for separate analysis. The aim is to examine whether ambiguity within selected Japanese architectural terms can be identified across the wider explanatory settings in which those terms are introduced, glossed, and discussed. Under this design, a term that shows contextual differentiation in FULL but weakens in CLEAN may indicate dependence on definitional support for its observed pattern. A term with similar differentiation across both conditions provides stronger evidence that its contextual pattern persists after definitional sentences are removed; representative uses must still be examined before interpreting that pattern as semantic ambiguity. The scholarly works cited in Section 4 serve as interpretive literature to contextualize the computational findings and are separate from the corpus evidence used in the analysis. These interpretive sources were not included in the FULL or CLEAN corpora used for the study-specific model fine-tuning or contextual analysis. The computational pipeline therefore had no direct study-specific exposure to these interpretive sources from which to reproduce their architectural readings. The representative contexts discussed in Section 4 are instead drawn from the corpus outputs generated by the computational pipeline and archived in the study repository. This separation is methodologically important because the scholarly literature provides an independent architectural frame for interpreting the computational results rather than evidence used to generate them.

2.3. Preprocessing and Normalization

The corpus text was normalized using a predefined mapping of orthographic variants and then divided into non-empty sentence units using the sentence-splitting rule specified in the configuration files. Candidate contexts were then identified by matching each target term within the normalized sentence set using Unicode-aware search. These sentence-level units provided the basis for the downstream ambiguity analysis, while the later extraction stage used target-centered span recovery to produce the sentence-level target vectors that entered clustering.

2.4. Active Model, Runtime Scope, Libraries, and Reproducibility Controls

This subsection retains technical detail because computational reproducibility is part of the study's methodological contribution. The purpose is not to treat model settings as

architectural findings, but to document the active computational environment and analytical configuration sufficiently for independent inspection and reproduction of the reported run. Figure 1 summarizes the broader analytical logic of these settings, and the following subsections develop it step by step. In practical terms, the workflow represents each occurrence of a target term from its surrounding sentence, groups similar contextual uses, compares the recovered patterns with broad architectural registers, and then tests whether those patterns remain visible after explicit definitional material is reduced. The technical settings below document how we executed these steps reproducibly.

The reported run used only the fine-tuned multilingual cased BERT masked language model. The runtime enabled contextual cluster identification for FULL and CLEAN, zero-shot architectural register alignment, weighted cluster-aware interpretation, report writing, and the focused engawa concept-alignment module. The engawa module used concept bootstrap, meaning that the focused hypothesis was evaluated against stored contextual occurrence embeddings. Zero-shot bootstrap and global visualization were inactive in the reported run.

The fine-tuned BERT model loaded for the reported run inherited the shared fine-tuning settings recorded in the YAML: learning rate 3×10^{-5}, per-device batch size 32, three training epochs, warm-up ratio 0.06, weight decay 0.01, maximum sequence length 128, masked-language-modeling probability 0.15, and gradient accumulation steps 2. The reported run used the fine-tuned model produced earlier in the training workflow.

The active pipeline used PyTorch 2.2.2 + cu121, Transformers 4.51.3, NumPy 1.26.4, pandas 2.2.2, scikit-learn 1.4.2, SciPy 1.13.1, regex 2025.9.1, PyYAML 6.0.2, and matplotlib 3.10.6. The codebase also included seaborn 0.13.2, networkx 3.3, gensim 4.3.3, and peft 0.17.1 for optional plotting, graph, Word2Vec, and adapter-loading branches, although the Word2Vec and base-model comparison paths were inactive in the reported run.

To improve reproducibility, the reported run was executed with fixed randomization settings and deterministic computation controls. The script fixed the random seed at 42, enabled deterministic PyTorch algorithms, disabled cuDNN benchmarking, and disabled TF32 fast paths. We preserved runtime logs, copied configuration files, the run manifest, the checkpoint-audit file, and exported figures as part of the reproducibility record for the reported run.

2.5. Screened Term Set

The active screening pool contained eight Japanese architectural terms: engawa, tokonoma, ma, mu, shakkei, doma, shōji, and fusuma. These eight terms form the operative set implemented in the script and the YAML file. The pool was selected because each term raises a plausible ambiguity question under the FULL and CLEAN conditions, while the set is small enough for a focused ambiguity study. Appendix A provides brief dictionary-based definitions as baseline lexical references for the eight target terms.

The eight-term set was purposively selected from the 28-term vocabulary examined in the first two studies of this research series; it was not intended to be exhaustive or statistically representative of that wider vocabulary. The terms were selected before the present computational analysis because prior architectural interpretation and literature review identified distinct ambiguity questions across materially anchored elements, boundary or transitional conditions, perceptual and representational relations, and more abstract conceptual cases. The set was kept deliberately small to allow detailed term-level contextual and architectural interpretation. Terms outside this set were not excluded because they were assumed to be semantically unambiguous; they fall outside the bounded analytical scope of the present study. Selection therefore focused on the architectural questions

posed by the terms, not on the stability, clustering, or evidential strength produced by the analysis.

The rationale varied across the eight terms. Engawa raised an ambiguity question around connective movement and boundary-making. Tokonoma was selected because it names a physically fixed interior site whose meaning can vary through display, ritual, and representation; the pre-analytical question was whether a materially stable alcove can also function as a representational site. Shakkei was selected because it complicates the relation between physical boundary and perceptual belonging: external scenery can remain physically outside a designed site while becoming compositionally incorporated into the garden [15,16]. Ma was included as a relational and temporal case spanning interval, spacing, pause, timing, and relation. Mu was included as a more abstract conceptual limit case centered on nothingness, absence, and potentiality, with less direct anchoring to a locatable architectural element. Doma brought the problem of an interior threshold whose material condition and use unsettle a simple inside-outside distinction. Shōji foregrounded a dividing surface that remains materially present while transmitting light and shifting between closure and connection. Fusuma raised an ambiguity question around movable partitioning, spatial recomposition, opaque separation, and possible image-bearing representation. The possible representational role of fusuma through fusuma-e was treated as a pre-analytical selection hypothesis informed by Stoichita's account of framed and embedded pictorial space that can stage fictive visual fields within a viewing situation [24]. The main pipeline examines whether fusuma contexts and register profiles support readings of partitioning, spatial recomposition, opaque separation, decorated surface use, or represented/imagined alignment. These pre-analytical interpretations motivated the selection of the target terms and are not derived from the present computational results. Section 4 develops their fuller architectural interpretations and examines them against the computational results, representative corpus contexts, and relevant scholarship.

We used the same main screening, clustering, and register-alignment procedure to analyze all eight terms. The only exception was engawa, which received an additional focused concept-alignment analysis as a secondary test of the boundary-and-connection hypothesis. No term-level manual priority, architecture-note, or override blocks were active in the reported YAML configuration, so the main screening, clustering, and classification logic treated the eight target terms uniformly.

2.6. Context Extraction and Contextual Clustering

Contextual cluster identification was the first substantive analytical stage after screening. For each target term, sentence contexts were retrieved by Unicode-aware whole-word regex matching after normalization. We then used contextual extraction for target-centered cropping before BERT tokenization, with a configured context window of 1200 characters and a maximum BERT input length of 256 tokens. This procedure prevented valid target occurrences from being lost when they appeared late in a long sentence. We used the fast tokenizer's offset mapping to identify the target token span within the cropped sentence.

The unit of analysis was the sentence-level context containing the target term. If a sentence contained multiple target occurrences, each occurrence vector was extracted and then averaged into one sentence-level target vector. Sentence-level fallback was turned off during contextual extraction, so no fallback vectors entered clustering; the fallback branch remained only as an extraction-audit field. Accordingly, matched-context counts denote the number of usable sentence-level target-span vectors entering downstream clustering, not the total number of raw mentions; the screening results are reported in Section 3.2.

We implemented contextual cluster identification as a conservative two-cluster probe of contextual variation that could warrant further interpretation as ambiguity. The active YAML capped automatic clustering at k = 2, with a minimum of ten matched contexts per corpus condition and a minimum cluster size of three. The resulting k value marks a two-way contextual contrast, not an exhaustive inventory of dictionary senses. For each accepted solution, three representative sentences from each cluster were stored for later interpretation.

K-means clustering used a fixed run seed along with the active YAML settings kmeans_n_init = 50 and kmeans_max_iter = 1000. Where non-trivial clustering was accepted and statistical diagnostics were enabled, we estimated stability via bootstrapped resampling and summarized it as the median adjusted Rand index. We assessed cluster separation via permutation testing and recorded it as an audit diagnostic. In this study, identifiable structure in the embedding space is called contextual differentiation. We use semantic ambiguity or polysemy as an interpretation of that structure only when representative contexts support substantively different meanings or functions; clustering metrics alone do not establish distinct senses.

2.7. Architectural Register Alignment

We performed architectural register alignment after contextual cluster identification in the pipeline, and kept the main interpretive frame deliberately tight: physical–spatial, symbolic–theoretical, represented/imagined, boundary/transition, and hybrid. These registers were used as limited analytical categories for comparing the target terms under FULL and CLEAN conditions, not as a complete ontology of architectural meaning. The five registers were selected as broad architecture-facing dimensions that can be applied across the heterogeneous target set without requiring the same specific architectural opposition to characterize every term. Physical–spatial captures material, locational, and built-spatial reference; symbolic–theoretical captures conceptual and abstract architectural discourse; represented/imagined captures representational, perceptual, or image-related meaning; boundary/transition captures threshold, edge, mediation, and spatial transition; and hybrid is reserved for cases in which multiple primary dimensions operate together without one adequately describing the identified profile. We did not use more specific architectural relations, such as inside/outside, light/dark, or artificial/natural, as universal primary registers because they are not equally applicable to all eight terms and would risk imposing term-specific architectural themes on the full target set in advance. Where relevant, we instead retained these relations in the representative contexts and secondary relation cues. We introduced the identified cluster centroids into the analysis during the later cluster-aware stage. Alongside these five primary registers, the active classification surface also included six secondary relation cues: boundary or edge conditions, threshold or transition, liminality, permeability or weak enclosure, display or ritual framing, and visual incorporation or borrowed scenery.

In implementation, the classification layer used eleven architectural prompt statements encoded in the same model space as the target-term representations. This alignment stage used target-term and prompt representations differently from the contextual clustering stage. The main clustering stage used sentence-level contextual target-span vectors extracted from corpus contexts. The raw register-alignment layer used standalone target strings and short architectural prompt statements as comparison anchors in the same fine-tuned BERT model space. These anchor-based similarities functioned as register-alignment diagnostics, while contextual target-span vectors remained the basis for clustering. The later cluster-aware stage reintroduced the identified contextual cluster centroids, so the final interpretation was not based only on undifferentiated standalone term repre-

sentations. Inference applied row-wise percentile thresholding at the 95th percentile together with a minimum cosine-similarity floor of 0.30, and no top-k override was used. Although the codebase included settings for consensus-style aggregation, only the fine-tuned multilingual BERT pathway was active in the reported run. The resulting register labels are conditional on the chosen architectural prompts, thresholds, and single model stream; they are not independently discovered semantic categories or cross-model consensus classifications. The hybrid register was one of the five primary registers, but it was not used as a residual category for unclear cases. We applied it only when the identified term-level or cluster-level profile combined multiple registers in a way that a single primary register could not adequately describe. Because hybrid-register use depended on interpretation, we checked each case against representative contexts before discussing it.

2.8. Cluster-Aware Analysis Procedure

The analysis included a cluster-aware stage after the initial register-alignment pass to preserve internal differentiation within each term. Cluster centroids identified during the contextual clustering stage were reloaded together with their cluster sizes. Each centroid was then compared against the same architectural prompt vectors used in the earlier register-alignment stage, and term-level profiles were recalculated from the identified cluster centroids using cluster-size-weighted aggregation. The cluster-aware stage reduced reliance on a single undifferentiated term average.

2.9. Focused Concept-Alignment Module for Engawa

The study applied a dedicated secondary concept-alignment module only to engawa. This module examined the specific architectural hypothesis that engawa may combine connector-flow and boundary-edge meanings, with the riverbank probe serving as a secondary analogy rather than an etymological claim. The configuration defined three concept axes, boundary_edge_axis, connector_flow_axis, and river_bank_probe, and limited concept alignment to engawa as the sole target. The script computed the derived hybridity measure from the boundary-edge and connector-flow axes instead of reading it directly from the YAML file. The study applied no equivalent bespoke concept-alignment module to tokonoma, shakkei, doma, shōji, fusuma, or any other term in the screened set. In this module, the raw concept-alignment table used standalone target and concept-anchor strings as focused comparison probes, while the concept-bootstrap stage evaluated the hypothesis against stored contextual occurrence embeddings. The module measured similarity between engawa and preselected conceptual probes. These exploratory comparisons do not establish that the probes are unique to engawa, outperform alternative probes, or provide evidence of etymology or historical origin.

2.10. Screening Thresholds and Retention Rules

Inclusion in the eight-term pool did not automatically make a term a strong ambiguity case. Under the active configuration, the minimum context requirement was ten matched contexts in FULL and ten in CLEAN. The minimum acceptable span-match rate was 0.60. A term met the stability criterion for a two-cluster contextual contrast only when its stability score reached the minimum threshold of 0.40. We reported lower-stability $k = 2$ solutions diagnostically when architecturally informative, but did not treat them as strong evidence of ambiguity. Persistence of contextual structure in CLEAN was required for a stronger cross-condition claim, but it was not sufficient by itself to establish semantic ambiguity. The screening configuration used a maximum fallback-ratio threshold of 0.40, although the reported run turned off sentence-level fallback during contextual extraction. The permitted shift in selected k between FULL and CLEAN was limited to one cluster. Cluster-separation permutation values were recorded as audit diagnostics. Terms with in-

sufficient context volume, low span identification, low stability, or weak cross-condition support were preserved only with greater interpretive caution. The screening stage served as an explicit filter within the analytical pipeline.

3. Results

3.1. Analytical Workflow

The results follow the pipeline order, with screening and contextual clustering setting limits on the subsequent architectural interpretation. The screening audit first establishes the support level for each target term under FULL and CLEAN corpus conditions. Contextual cluster identification then examines recurring internal structure. Architectural register alignment translates that structure into five interpretive registers: physical–spatial, symbolic–theoretical, represented/imagined, boundary/transition, and hybrid. Cluster-aware interpretation then compares the identified contextual clusters within each target term, preserving internal differentiation. A final focused engawa module serves as a secondary test of one architectural hypothesis and is not repeated across the full target set. The workflow summarized earlier in Figure 1 gives the procedural order through which the evidence accumulates.

Context volume, cluster stability, representative uses, and prompt alignment jointly limit the strength of each architectural reading. The section begins with recurrent contextual structure, then moves to architectural register alignment and interpretation. The target terms play different roles within this sequence. For interpretation, the terms show different evidence profiles: relatively stable contextual structure in several cases; broad or abstract patterns in ma and mu; a cautious case in fusuma; and a low-stability diagnostic case in doma. These differences matter because the strength and semantic interpretation of contextual structure vary across the target set. This section uses representative sentence cues sparingly to show that the identified clusters correspond to recognizable contextual uses. The Section 4 provides the fuller architectural interpretation of these representative contexts.

3.2. Screening Audit Under FULL and CLEAN Conditions

Table 1 summarizes the initial support level for each target term under the paired corpus conditions. Matched-context volume refers to usable exact-span matches that enter downstream clustering, not raw occurrence totals. Context volume, identified cluster number, span identification, and stability must therefore be considered together when assessing each term's evidential strength.

Mu and fusuma remain usable but narrower. Mu is stable across both corpus versions, but it is less directly anchored to a locatable architectural element than engawa, tokonoma, shakkei, shōji, or fusuma. Fusuma also retains a two-cluster solution, but weak separation diagnostics make its interpretation more cautious; its evidence supports movable partitioning and spatial recomposition more strongly than image-bearing representational extension. Doma should be treated only as a low-stability diagnostic case. Although it meets the minimum context and span requirements, its stability is too low to support a strong ambiguity claim. Overall, the screening audit shows uneven support across the eight terms, with ambiguity claims scaled to the strength of evidence for each term.

Table 1. Screening audit under FULL and CLEAN conditions, reporting usable matched-context counts, identified cluster number, span identification, and stability across the target set.

Terms	Full Matched Contexts (*n*)	Clean Matched Contexts (*n*)	Full k	Clean k	Full Span	Clean Span	Full Stab	Clean Stab
Doma	20	20	2	2	1.00	1.00	0.22	0.22
Engawa	188	155	2	2	1.00	1.00	1.00	0.92
Fusuma	49	37	2	2	1.00	1.00	1.00	1.00
Ma	572	535	2	2	1.00	1.00	0.98	1.00
Mu	51	49	2	2	1.00	1.00	0.92	0.92
Shakkei	94	83	2	2	1.00	1.00	1.00	1.00
Shōji	215	196	2	2	1.00	1.00	1.00	1.00
Tokonoma	86	79	2	2	1.00	1.00	1.00	1.00

Note: Matched contexts (*n*) refer to usable sentence-level target-span vectors entering downstream clustering, not raw mention totals. Span and stability are audit metrics, not semantic-accuracy scores. We retained cluster-separation permutation diagnostics in the screening decision outputs and discuss them where they affect interpretation. However, they are not reported in Table 1 to keep the screening audit readable. The values are available in the public repository, especially in screening_decision_table.csv, summary_dashboard.csv, and master_results_long.csv. A span value of 1.00 indicates complete target-span identification among candidate sentence contexts, while a stability value of 1.00 indicates the maximum observed median bootstrap ARI for the accepted k = 2 solution under the active run settings. The screening audit separates the eight terms into three evidential groups. Engawa, tokonoma, ma, shakkei, and shōji form the strongest group for recovery of contextual structure because they combine adequate context volume, complete span identification, stable two-cluster solutions, and persistence across both corpus versions. Ma has the largest context base, while engawa, shakkei, shōji, and tokonoma show especially strong stability profiles. Engawa and tokonoma still require small qualifications: engawa has weaker cluster separation in CLEAN, and tokonoma has weaker cluster separation in FULL. These limits affect the strength of those diagnostics, but not their overall retention as strong cases.

3.3. Contextual Cluster Identification Across Target Terms

Figure 2 shows recurrent internal structure across the target set. Bar length indicates matched-context count, while colored segments mark identified clusters aligned within each target term across FULL and CLEAN. The representative sentence cues serve only as brief orientation markers at this stage; fuller architectural interpretation is deferred until later.

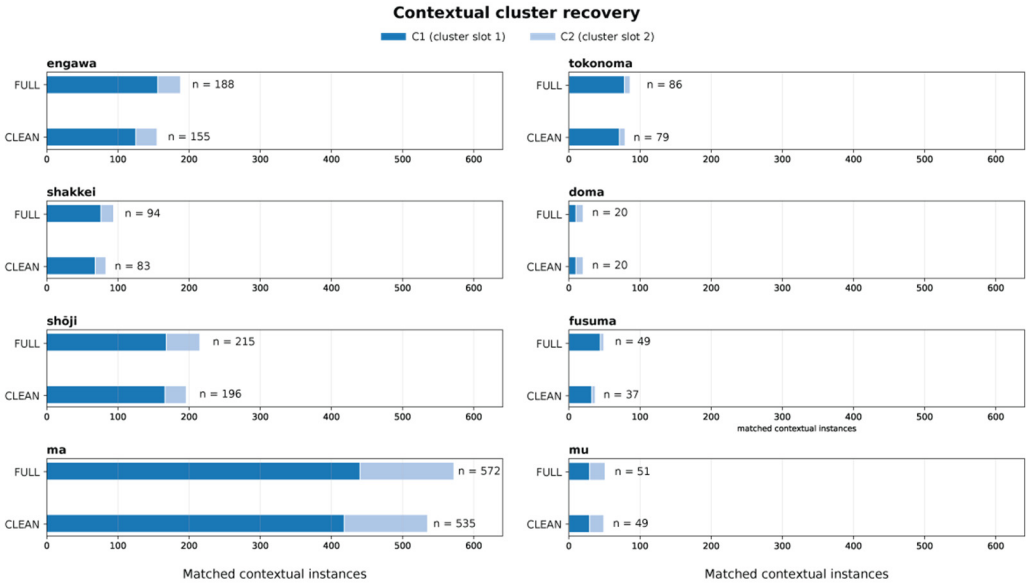


Figure 2. Contextual cluster identification across target terms under FULL and CLEAN conditions. Bar length indicates matched-context count, while colored segments represent identified contextual clusters aligned within each target term across corpus conditions.

Engawa showed a stable two-cluster structure in both corpus conditions. One recurrent usage mode treated engawa as a physical edge or extension, with contexts including roof evolution, the projecting edge of the house, ground slope, a traditional stone step, and drainage-related mediation between indoors and outdoors. The other usage mode treated engawa as a connective, mediating space, with contexts including an intermediate area between interior and exterior, a veranda, a connector between people and nature, and a privacy-related spatial system. FULL and CLEAN differed in balance, but both showed the same conservative binary structure. Tokonoma likewise displayed a clear two-cluster structure under both corpus conditions. One recurrent mode was anchored in the alcove as a display and room-orientation site, including object placement, flower arrangement, art display, host-guest positioning, and adjacency to a built-in writing table. The other was more closely tied to named tokonoma discourse, including short title-like contexts such as "Tokonoma", "Tokonoma at Tenryū-ji", and "Tokonoma: A Showcase of Ma". The tokonoma split appeared under FULL and CLEAN and had strong span identification and stability, although its title-like cluster offers limited evidence for a distinct meaning. The alcove-and-display cluster is more contextually informative than the title-like cluster.

Shakkei divides between scenic-exemplary uses and more concept-framing uses in both corpus conditions. One recurrent mode was definitional, exemplary, or scenic, including questions about how to define shakkei, photographic examples, subtle examples, and incorporating natural-environment elements into yard design. The other was more title-like or concept-framing in CLEAN, while in FULL it included shakkei as an ancient concept, its possible incorporation into tall Japanese structures such as pagodas, and the attempt to capture nature alive. Mu showed a smaller but still legible two-cluster structure, with one recurrent mode oriented toward philosophical, Buddhist, or architectural uses of nothingness, absence, negation, non-relativity, and the absence of substance, and another toward lexical, etymological, translation, or cross-linguistic glossing of lack, without, emptiness, void, wu, and "not". These cases differ in architectural anchoring, and their representative contexts show that the identified clusters correspond to recognizable patterns of use.

Ma had the largest matched-context base and showed a two-cluster structure. Its representative cues suggested a broad contrast between literal interval or between-space usage, including ai-no-ma and temporal "ma" in a story, and more explicitly conceptual language of space, pause, emotional expression, balance, harmony, simplicity, and meaning. Doma showed a plausible split between entrance, adjacency, raised-floor relation, and earthen-floor access on one side and service-circulation use on the other. Still, its low stability gives it less interpretive weight than the stronger cases. Shōji separated between sliding-panel construction, grooved tracks, wooden frames, portability, room-joining, and privacy on one side and named subtype or technical-variant contexts on the other. Fusuma showed two clusters in both corpora, with representative contexts centered on sliding doors, room separation, internal reconfiguration, opacity, decorated surfaces, paintings, and typical Japanese-room composition. The accepted $k = 2$ results show recurrent contextual contrasts under the configured binary probe, although neither the number of underlying senses nor the semantic significance of each contrast follows from k alone.

3.4. Architectural Register Alignment at Term Level

Figure 3 shows how each target term leans across the five main registers once recurrent internal structure has been established. Because the scores are rescaled within each term across FULL and CLEAN, the figure supports within-term comparisons more strongly than cross-term comparisons.

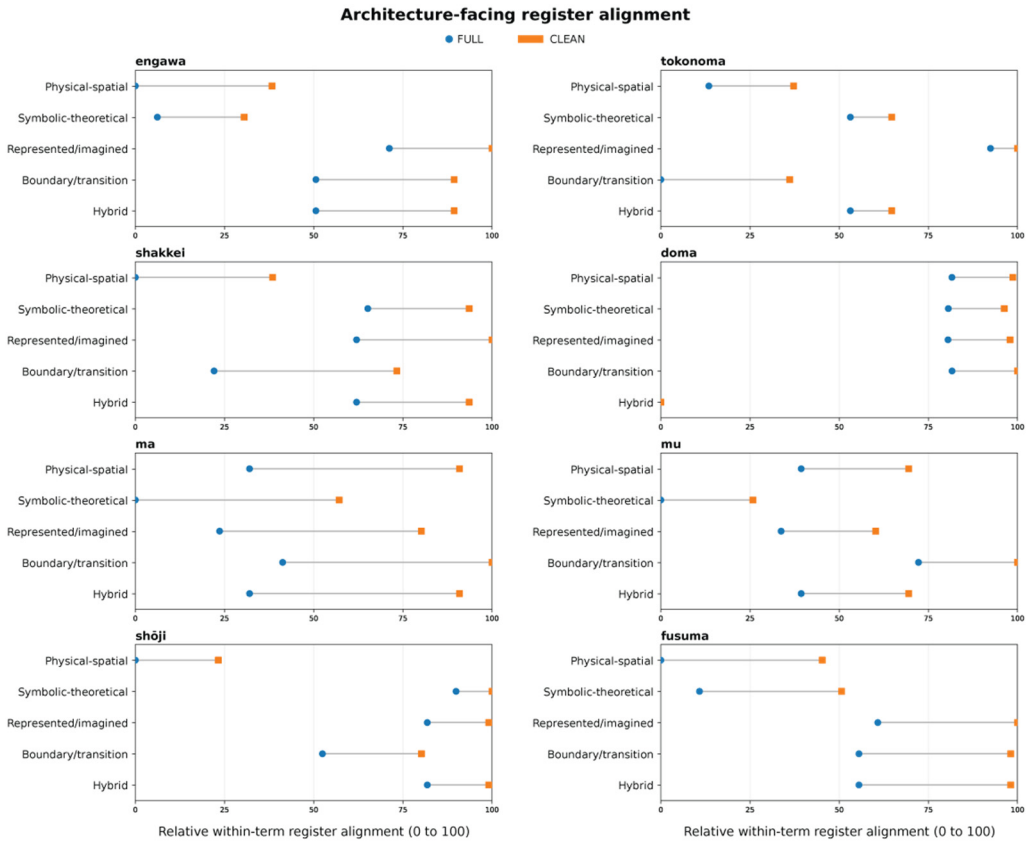


Figure 3. Architectural register alignment across target terms. Scores are rescaled within each term across FULL and CLEAN together and indicate relative within-term register alignment of the five primary interpretive registers.

Engawa showed the strongest alignment with the represented/imagined, boundary/transition, and hybrid registers, while its physical–spatial weighting was relatively weak. Engawa's alignment with the represented/imagined, boundary/transition, and hybrid registers strengthened under CLEAN. Tokonoma likewise resisted reduction to a merely physical–spatial reading. Its strongest signal was represented/imagined, supported chiefly by symbolic–theoretical and hybrid alignment, while boundary/transition was weaker and more condition-dependent. At the term level, engawa and tokonoma both refer to recognizable built elements, but their register profiles also include symbolic, representational, boundary-related, or hybrid meanings.

Shakkei produced one of the clearest profiles, with weak physical–spatial dominance and strong symbolic–theoretical, represented/imagined, and hybrid alignment. Boundary/transition was present but more condition-dependent, while physical–spatial weighting was weak in FULL and rose under CLEAN. Ma showed broad alignment across several registers, especially physical–spatial, represented/imagined, boundary/transition, and hybrid, with symbolic–theoretical alignment becoming more visible under CLEAN. Mu's contexts clustered reliably, but the clusters were tied more to philosophical and lexical uses than to a locatable architectural element. Shōji showed strong symbolic–theoretical, represented/imagined, and hybrid alignment, with boundary/transition present but not dominant. Doma and fusuma produced interpretable register profiles, although doma shows low stability and fusuma supports movable partitioning and recomposition more directly than representational extension.

3.5. Cluster-Aware Interpretation

Figure 4 brings together two levels of comparison. The left panels show the raw matched-context counts for each identified contextual cluster under FULL and CLEAN, while the right panels show the register composition within those clusters. This cluster-level comparison extends the analysis beyond cluster size alone, as the clusters can be interpreted as distinct yet related contextual configurations.

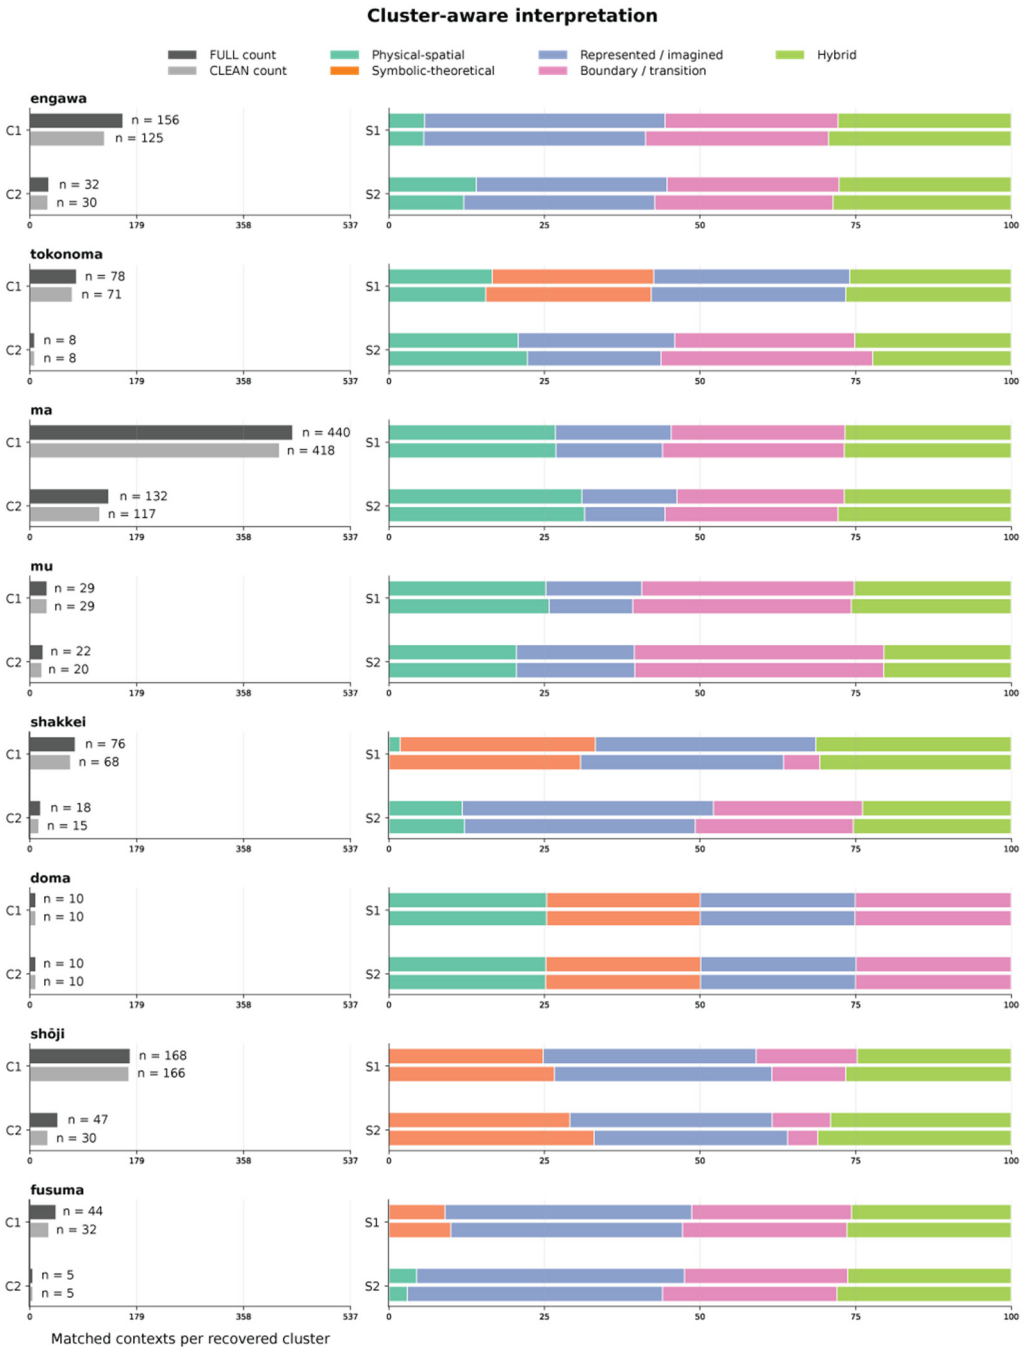


Figure 4. Cluster-aware interpretation across target terms. Left panels show matched-context counts for recovered contextual clusters under FULL and CLEAN. The right panels show the within-cluster register composition across the five primary interpretive registers. All five registers are evaluated for every cluster, but a register may have no visible segment when its contribution is not retained in the resulting cluster-level profile; the absence of a color therefore indicates no displayed contribution for that register, not that the register was omitted from the analysis.

Engawa produced one of the clearest cluster-aware results. One cluster was larger than the other in both FULL and CLEAN, but the CLEAN condition created no new dominance pattern. Both clusters were composed chiefly of represented/imagined, boundary/transition, and hybrid registers, instead of dividing into a simple physical–spatial-versus-abstract opposition. The shared register composition is important because the representative contexts suggested two related but distinct uses: one oriented toward edges, extensions, roofs, grounds, steps, summer use, and drainage mediation, and the other toward veranda-like connections, interior–exterior mediation, people-nature relations, and privacy. The cluster-aware comparison shows that both are architecturally layered. Engawa's clusters suggest structured contextual variation within a boundary-and-connection field; whether that variation constitutes distinct senses remains an architectural and semantic interpretation.

Tokonoma also produced a clear but modest cluster-level contrast. Both clusters appeared under FULL and CLEAN, and both had a strong component of representation/imagery. One cluster was more contextually anchored in built alcove use, object placement, flower arrangement, art display, guest orientation, host display, seated eye level, and adjacent furnishing. The other was more title-like or concept-framing, including "Tokonoma," "Tokonoma at Tenryū-ji," and "Tokonoma: A Showcase of Ma." Shakkei showed two clusters within a shared visual-conceptual field. The cluster-aware composition supports mainly visual, symbolic, and hybrid readings, but the two clusters have different register profiles. One cluster leans more clearly toward symbolic–theoretical and represented/imagined alignment, while the smaller cluster gives more weight to represented/imagined, boundary/transition, and hybrid composition. The result supports perceptual or scenic incorporation, although the two clusters have different register hierarchies. Mu showed a small but interpretable two-cluster structure, with one side leaning toward philosophical, Buddhist, and architectural uses of nothingness or absence, and the other toward lexical, translation, etymological, and cross-linguistic contexts.

At the cluster-aware stage, ma, doma, shōji, and fusuma separate into distinct support profiles. Ma's two clusters were compositionally close enough to support the interpretation of semantic breadth, not a cleanly defined opposition. Doma's two clusters were architecturally plausible, but the low stability shown in the screening audit makes the pattern diagnostic rather than strong evidence. Shōji showed a stable two-cluster structure, with one cluster oriented toward sliding-screen construction, grooved tracks, wooden frames, portability, room-joining, and privacy, and the other toward named subtypes or technical variants such as koshi-tsuki-shōji, kake-shōji, and ryōmen shōji. Fusuma showed two clusters in both corpus conditions. Still, the smaller second cluster, weak separation diagnostics, and representative contexts support a cautious interpretation: the term is identifiable as a movable partition, an opaque divider, a decorated surface, and a recompositional device, while the richer representational-extension reading is only suggestive.

3.6. Bounded Engawa Hypothesis-Stress Test

Figure 5 reports the focused engawa probe-alignment test. Higher cosine similarity indicates stronger alignment between engawa and the relevant conceptual axis. The figure makes no ontological or etymological claim; it examines whether one architectural interpretation has measurable alignment in the model space under FULL and CLEAN conditions.

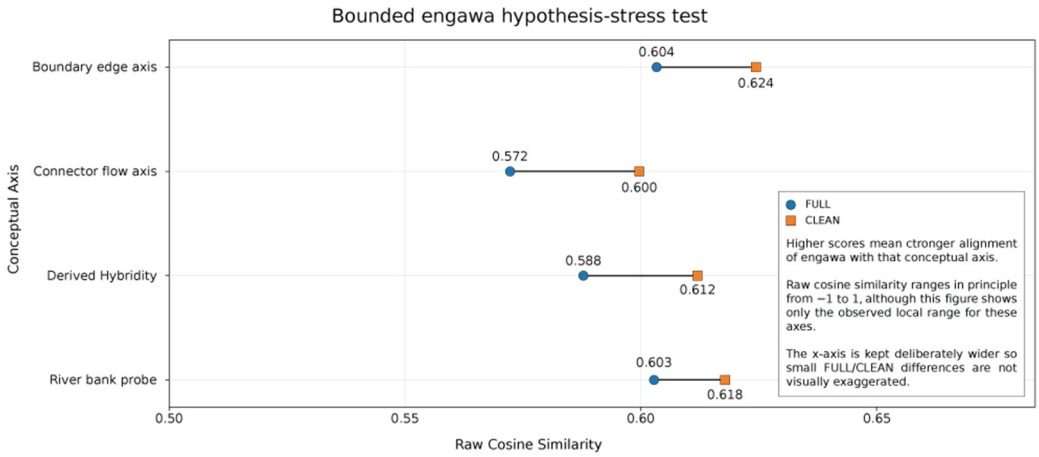


Figure 5. Bounded engawa hypothesis-stress test, presented as an exploratory concept-alignment probe. Cosine similarity scores compare engawa under FULL and CLEAN conditions against a boundary-edge axis, a connector-flow axis, a derived hybridity axis, and a riverbank analogy probe.

All four probes showed positive cosine similarity, and all four had slightly higher values under CLEAN than under FULL. The boundary-edge axis rose from 0.604 in FULL to 0.624 in CLEAN. The riverbank probe was similar, with scores of 0.603 and 0.618. The connector-flow axis was lower but still positive, at 0.572 and 0.600. Derived hybridity fell between these values at 0.588 and 0.612.

This pattern is consistent with a focused interpretation of engawa as both an edge condition and a connector. Boundary-edge and riverbank have the highest raw alignment values in the concept-alignment plot, while connector-flow also shows positive alignment in both corpus conditions. Within the selected probe set, the model gives positive similarities for boundary, connection, hybridity, and the riverbank analogy; those values do not establish that any one probe is distinctive or dominant. The riverbank probe is not evidence of etymology or historical origin. It indicates only that the riverbank analogy has measurable alignment within the corpus-model space used in this study.

3.7. Cross-Stage Synthesis

Taken together, the results show several distinct ambiguity outcomes across the target set. Engawa and tokonoma emerge as two central cases because both show internal structure under FULL and CLEAN and remain interpretable at both the term and cluster level without collapsing into simple material description. Tokonoma is supported mainly by its span-identification and stability results because one of its two clusters is title-like and less contextually rich. The engawa results further support a focused hypothesis of combined boundary-bearing and connective logic. Tokonoma shows that a physically located architectural site can be semantically layered, with one cluster tied to object placement, display, host-guest positioning, and adjacent furnishing, and the other functioning more as a title-like or concept-framing cluster. Shakkei also performs strongly across the general pipeline and is one of the clearest non-bespoke ambiguity cases in the set. Its results consistently point to a mainly represented/imagined, symbolic–theoretical, and hybrid ambiguity organized around scenic incorporation and uncertain spatial belonging.

Ma and shōji are strong but differently structured cases, while mu and fusuma require greater caution. Ma shows a broad semantic range around interval, relation, pause, timing, balance, harmony, simplicity, and meaning, not a narrow dictionary-like bifurcation. Mu is interpretable but modest, functioning as the more abstract comparison case in the set. Shōji shows a stable two-cluster structure and supports interpretations of sliding construction, paper-and-frame materiality, room-joining, privacy, named variants, lumi-

nous permeability, and flexible enclosure. However, its register profile is not dominated by boundary/transition alone. Fusuma supports partitioning, opacity, decorated surfaces, paintings, and spatial recomposition more clearly than a strong claim about representational extension. Doma is the lowest-stability case, and its low stability limits the force of its otherwise plausible interior-threshold pattern. The results support a qualified conclusion: under audited FULL and CLEAN corpus conditions, the contextual structure associated with ambiguity may persist, broaden, sharpen, or weaken depending on the term and the discourse condition.

4. Discussion

Architectural scholarship treats vocabulary as conceptually active and historically unstable [1], while work on Japanese architecture foregrounds relations, intervals, rituals, domestic organization, and ambiguity beyond isolated objecthood [9,12]. Japanese aesthetics extends this frame: Marra links meaning to language, representation, perception, nature, and structures of significance [13]; Saito connects aesthetic experience to ordinary life and ethical orientation [14]; and Amit reads Japanese aesthetic terms relationally within a wider conceptual structure [6]. Against this background, the present study tests whether ambiguity developed qualitatively in prior scholarship remains identifiable in a broad expository corpus after overt definitional support is reduced.

Across the screening audit, contextual cluster identification, architectural register alignment, and cluster-aware interpretation, the eight terms separate into several profiles: comparatively stable contextual differentiation, broad semantic-field structure, modest but legible abstraction, cautious partitioning, and low-stability diagnostic evidence. Because all eight terms passed through the same main pipeline, these differences are analytically informative. The computational pipeline identifies recurring contextual structures and register alignments but does not independently establish architectural concepts or final semantic categories. The architectural formulations developed below therefore interpret those patterns through representative corpus contexts and existing scholarship. The representative contexts come from pipeline-generated CSV files archived in the study's public repository and provide the textual basis for interpreting the cluster structure. Each subsection accordingly distinguishes the computational evidence reported in Section 3 from the architectural interpretation developed from that evidence.

4.1. Engawa as a Boundary-Making Mediator

Engawa is one of the strongest cases in the study because its contextual structure persists across FULL and CLEAN and is additionally examined through the focused concept-alignment test, although the CLEAN permutation-separation diagnostic remains weaker. The representative contexts distinguish two related tendencies: one centers on physical edge, extension, ground/step conditions, and drainage mediation, while the other centers on veranda-like mediation, privacy, people-nature relations, and interior–exterior connection. The architectural significance of this result lies not in treating these tendencies as unrelated senses, but in examining how boundary production and connection operate through the same spatial condition. Existing scholarship approaches the mediating dimension of engawa from several directions. Wong [18] frames engawa through ring-like circulation and inside-outside transition, while Hou et al. [19] describe it as a dynamic spatial mediator shaped by boundary permeability and sequential flow. These direct engawa studies place the term beyond the category of a veranda-like appendage or a vague in-between zone. Yang's [20] work on shikii-related threshold spaces provides a comparative frame for graduated boundaries, especially where domestic, commercial, and social interactions

meet across layered spatial edges, but it should be read as a parallel threshold framework rather than as direct evidence of engawa.

Read alongside these studies, the engawa results suggest that boundary production warrants attention alongside mediation. Engawa connects the interior and exterior while also preserving their differences. It softens the boundary between house and garden without erasing it. Its ambiguity depends on this double operation: the same spatial element produces both connection and separation.

A riverbank analogy helps clarify this double operation in two directions. Along its course, a river connects spatially separated places through continuous flow; across its course, the same river distinguishes two opposing banks while keeping them in relation. Engawa can be read similarly: it supports circulation and connection along the building edge while, across that edge, marking and mediating the distinction between interior and exterior. The focused engawa test aligns positively with boundary-edge, connector-flow, derived hybridity, and the riverbank analogy, although the bootstrap differences are too weak to establish one axis as dominant. The riverbank relation is therefore used as an architectural analogy rather than as an etymological or historical claim. The selected computational probes are compatible with the boundary and connector components of this interpretation; the longitudinal and transverse spatial reading is an architectural formulation developed from the corpus contexts and literature.

Kurokawa's account frames engawa as a third type of space, neither simply interior nor exterior, and this remains the closest architectural precedent for the present interpretation [17]. Based on the computational evidence reported above and this architectural precedent, the paper proposes modal littoral-circuit space as an interpretive formulation rather than as a semantic category established directly by the model. The formulation specifies one way of reading how this thirdness operates: engawa connects house and garden by maintaining, thickening, and inhabiting the edge between them. "Modal" refers to its changing spatial status: the same engawa may function as an exterior-facing edge, an interior extension, a sheltered passage, an informal sitting zone, a privacy buffer, or a domestic circulation strip, depending on enclosure, exposure, season, weather, use, and movement. In traditional houses, including minka, this condition could also be physically altered by using sliding amado at the outer edge of the engawa. When the shutters were retracted, the engawa remained open toward the exterior; when drawn across the outer edge, particularly at night or during adverse weather, they enclosed the veranda and produced a more interior-like spatial condition [44,45]. The spatial status of engawa was therefore not determined by location alone but could change with enclosure, environmental conditions, and occupation. "Littoral" names the edge-zone condition: like a shore, it holds two domains in contact without reducing them to a single domain. "Circuit" identifies its role in movement along the building perimeter, where circulation occurs through the inhabited edge as much as through direct crossing between inside and outside. This formulation offers one architectural reading of engawa as an intermediary edge condition shaped by enclosure, boundary production, movement, and domestic use.

4.2. Tokonoma as Symbolic Interiority

Within the reception room, tokonoma combines a fixed architectural form with changing display, social, ritual, and representational roles. The clusters do not form a simple opposition between physical alcove and abstract meaning. One mode is materially anchored in the alcove as a display and room-orientation site. In the representative contexts, tokonoma appears next to a built-in writing table, as an alcove where paintings, pottery, flower arrangements, and other art forms are displayed, as a site where object placement is considered from the eye level of a person seated on tatami, and as a setting in which

the flower arrangement indicates the relative wealth of the host even though the guest sits with their back to it. The second mode is more title-like or concept-framing, including "Tokonoma", "Tokonoma at Tenryū-ji", and "Tokonoma: A Showcase of Ma". Together, these contexts support tokonoma as a room-facing site where display, viewing position, social orientation, furnishing, and aesthetic concentration are joined.

The architectural interpretation developed from these results concerns tokonoma's representational capacity within a fixed interior site. Scholarship on tea ritual and symbolic placement treats objects, orientation, and spatial arrangement as part of the meaning of the setting [21–23]. Hojda describes tokonoma as a concentrated site for representing nature, in which seasonal flowers and hanging scrolls make natural and seasonal associations present within the interior [46]. The Metropolitan Museum of Art likewise documents flower arrangements and calligraphic or painted scrolls displayed in tokonoma, including Zen calligraphy valued for spiritual and didactic qualities [47]. Pilgrim's account of Japanese art as a religio-aesthetic field further supports the view that aesthetic form and symbolic presentation can carry spiritual and conceptual significance [48].

The computational results can then be compared with the pre-analytical interpretation that motivated tokonoma's selection. Because the recovered contrast lies between the two contextual modes described above, it does not establish a discrete opposition between a physical interior alcove and a represented exterior or abstract conceptual space. Read together with the representative contexts and the cited scholarship, the results support interpreting tokonoma as a physically stable interior site whose changing contents can make domains beyond its material location present through representation. Ikebana can bring nature and seasonal change into the room; landscape painting can introduce an imagined spatial field; and poetic or Zen calligraphy can introduce literary, philosophical, or spiritual meaning. The alcove does not cease to be a physical space or literally become an abstract one. Its ambiguity lies in the simultaneous operation of material location and representational extension: the stable alcove remains part of the interior and continues to organize display, ritual, and social orientation, while its displayed contents make natural, spatial, and conceptual domains present within it. The clustering does not independently establish this full representational interpretation; it provides the contextual structure from which the broader architectural reading is developed together with the representative contexts and cited scholarship.

4.3. Shakkei and the Ambiguity of Incorporation

Shakkei also performs well in the main pipeline, though it lacks a secondary concept-alignment test comparable to the engawa module. It has two clusters across FULL and CLEAN, with complete target-span identification and maximum observed bootstrap stability. Its architectural register profile is among the clearest non-physical–spatial profiles in the set. Its strongest alignment lies in the symbolic–theoretical, represented/imagined, and hybrid registers, while boundary/transition is secondary. The interpretation turns on how external scenery becomes compositionally incorporated into the designed field. Shakkei, usually translated as borrowed scenery, is more than a view from a garden. Representative contexts include definitional questions, photographic examples, subtle scenic examples, the incorporation of natural-environment elements into yard design, ancient-concept framing, design-principle framing, and the attempt to capture nature alive. The ambiguity concerns scenic incorporation: an external or surrounding landscape can stay physically outside the designed site while entering its perceptual and compositional field.

Existing literature on borrowed scenery already treats shakkei as a compositional relation in which external scenery becomes part of the garden's designed field [15,16].

The present results provide corpus-based support for that qualitative account, but the identified split still requires caution. In CLEAN, one cluster includes definitional and exemplary uses, while the other includes more title-like or design-principle contexts. In FULL, the contrast is richer, extending from natural-environment incorporation and scenic capture to ancient-concept and tall-structure examples such as pagodas. The term's ambiguity lies in how scenery, background, photograph, natural environment, yard design, and surrounding landscape become compositionally available within the designed field. Taken together, the corpus evidence and the qualitative literature [15,16] support interpreting shakkei as perceptual incorporation and ambiguous spatial belonging. The ambiguity lies in the mismatch between physical location and compositional membership: the borrowed landscape remains outside the site's material boundary, yet through visual framing and composition it becomes part of the garden as perceived and designed.

4.4. Ma and Mu: Breadth, Abstraction, and Controlled Multiplicity

Unlike the more sharply contrasted cases, ma spans overlapping spatial, temporal, experiential, and conceptual uses; its representative contexts support relational breadth without establishing two discrete senses. Ma extends across literal interval, the space between structures, named between-spaces such as ai-no-ma, temporal timing in speech or story, spaces or pauses in expression, emotional communication without speech, and aesthetic language of balance, harmony, meaning, simplicity, and the exclusion of unnecessary elements. Its uses overlap across spatial, temporal, experiential, and conceptual registers.

Reading ma as relational semantic breadth accords especially with Pilgrim's account of ma as a spatial and temporal interval, pause, opening, and relation [10]. More broadly, Marra [13] supports the adjacent interpretive point that Japanese aesthetic meaning can resist stable object reference and operate through layered or deferred signification.

The diffuse two-cluster pattern is methodologically important because semantic richness can emerge without sharply separated clusters. Some terms support ambiguity by sustaining distinct contextual roles; others are semantically dense because their uses span neighboring registers. Ma belongs closer to broad semantic spread than to a clean two-sense opposition. Its recurrent contextual organization remains diffuse, and its architectural significance lies in relational spacing rather than a bounded object or a single conceptual operation.

Mu is more modest but useful as an abstract comparison case. It has two clusters across FULL and CLEAN, with complete target-span identification and stable clustering in both corpus conditions. One mode leans toward philosophical, Buddhist, or architectural formulations of nothingness, absence, negation, non-relativity, and the absence of substance, including the Benesse House Oval example and the relation between absence, boundaries, and viewing capacity. Another leans toward lexical, etymological, translation, or cross-linguistic glosses such as emptiness, void, lack, without, wu, and not. Unlike engawa, tokonoma, or shakkei, mu is weakly tied to a locatable architectural element. Mu's ambiguity is more conceptual than spatial. Mu's value is comparative: it marks the abstract end of the screened field and shows that the pipeline can identify internal multiplicity even where architectural anchoring is weaker.

The earlier pairwise analysis in the same research program provides a complementary comparative context for these within-term results [43]. That analysis found temporal pause, compositional balance, and negative-space fields to be more consistently ma-directed, whereas mu-facing nothingness, nonbeing, and potentiality were directionally coherent but less consistently supported. The present analysis addresses a different question: ma shows internal breadth across relational, spatial, temporal, and aesthetic uses, while mu shows a more abstract internal contrast between philosophical or architectural

formulations and lexical or translational uses. The two analyses should therefore be read as related analytical perspectives rather than as independent validation of one another.

4.5. Doma as an Interiorized Threshold

Doma has a plausible architectural interpretation, but its two-cluster solution has low stability in both corpus conditions. Doma complicates the inside-outside relation from within the house itself. A doma belongs to the dwelling's interior organization, yet differs from the raised, clean interior. Its earth-floored or unfloored condition, association with entry, service/circulation, adjacency, partitioning, and movement from the front of the house to the rear give it an exterior-like status inside the house.

Doma's ambiguity comes from the overlap between exterior-like materiality and interior organization. Doma can be read as an interiorized threshold: a zone where exterior materiality and domestic interiority overlap. It manages the transition before entry into the raised living space is fully completed. It is enclosed, yet it carries traces of the outside. It belongs to the house while differing from the most socially or materially protected interior zones. The representative contexts support this reading by describing a main entrance providing access to the doma, an earth-floored area that provides movement from the front of the house to the rear, and an internal division between unfloored service/circulation space and raised tatami- or timber-floored rooms.

The limitation of the doma result concerns evidential strength rather than architectural plausibility. With only 20 matched contexts in each corpus condition and low stability under both FULL and CLEAN, the interiorized-threshold reading should remain suggestive rather than being treated as a strong ambiguity case.

4.6. Shōji as Luminous Boundary

Shōji requires careful interpretation because the corpus results support construction, permeability, and flexible enclosure more directly than a simple boundary/transition claim. Shōji refers to a partition that separates space without functioning as a fully closed wall. A shōji separates rooms or zones while also transmitting light. A shōji blocks direct vision without producing opaque closure. It can be opened, removed, or slid aside, so its boundary function can shift between separation and connection [25].

The architectural interpretation developed here is that shōji functions as a light-filtering, adjustable partition, grounded in the identified constructional and functional contexts. Its ambiguity involves division and connection, but the representative contexts point more concretely to sliding panels, grooved wooden tracks, paper attached to thin wooden frames, panels that can be pushed open to join rooms and closed for privacy, portable or adjustable construction, and named subtypes such as koshi-tsuki-shōji, kake-shōji, and ryō-men shōji. The luminous-boundary reading is therefore treated as an architectural interpretation drawn from constructional, material, and functional cues in the identified clusters.

This interpretation is consistent with accounts of Japanese architecture and domestic space that discuss sliding screens, translucent materials, flexible spatial organization, and the negotiation of boundaries between interior and exterior zones [49], as well as private and public domestic zones [25]. The stable two-cluster recovery supports reading shōji as both a material partition and a flexible boundary device.

Shōji has a stable contextual split, although one cluster emphasizes named technical variants, and its theory-guided register profile is led by symbolic–theoretical, represented/imagined, and hybrid alignment, with boundary/transition remaining secondary. The luminous-boundary formulation should therefore be read as an architectural interpretation of the constructional and functional evidence, not as a claim that boundary/transition dominates the register signal.

Within the broad expository corpus used here, the analysis supports shōji more directly as a light-transmitting, permeable, and adjustable enclosure than as a device of perceptual spatial enlargement. Architecturally, accounts of shōji as translucent, light-filtering partitions that visually connect spaces and participate in flexible boundary conditions support a further interpretation [25,49]: light passing through a closed shōji may make the enclosure feel less final than an opaque partition and suggest continuity with the space beyond, even though the physical boundary remains in place. The present analysis does not directly examine this proposed perceptual effect, which remains a secondary architectural interpretation, and not a computational finding.

4.7. Fusuma and the Limits of Representational Extension

Fusuma is a cautious case because the evidence supports partitioning and spatial recomposition more strongly than representational extension. The representative contexts point most directly to sliding doors, partitions, room division, opacity, decorated interiors, paintings, typical Japanese-room composition, and flexible internal configuration. These contexts support several recurrent uses of fusuma, especially partitioning and room reconfiguration, while the proposed representational extension remains a more tentative architectural interpretation.

The representational-threshold reading of fusuma is this paper's architectural interpretation, not a direct claim drawn from Stoichita. Stoichita's account of framed and embedded pictorial space helps explain how pictorial devices can create or stage fictive visual fields within a viewing situation [24]. By analogy, a painted fusuma may remain an active architectural partition while also carrying a pictorial field. When fusuma-e presents landscape, scenery, figures, or other spatial imagery, the represented field does not physically enlarge the room. Still, it can alter the perceived limit of an otherwise confined interior by introducing pictorial depth or an imagined domain beyond the opaque partition surface. In this sense, fusuma-e can produce a perceptual or representational extension through the boundary itself. This operation relates to, but is architecturally distinct from, the interpretation of tokonoma: tokonoma stages another natural, symbolic, or conceptual domain within a fixed alcove, whereas fusuma-e turns the enclosing partition surface itself into the represented field. This interpretation remains compatible with accounts of Japanese domestic interiors in which fusuma panels mediate spatial and social relations between domestic zones [25]. Yamamoto further supports the image-bearing possibility by showing that fusuma paintings can surround rooms, differentiate interior atmospheres, and become incorporated into the architecture of a room or building [50]. Computationally, fusuma is not a weak case: it retains a two-cluster solution under both FULL and CLEAN, with complete target-span identification and maximum observed bootstrap stability in both conditions. The caution concerns interpreting that structure rather than recovering it. The representative contexts directly support movable partitioning, opaque separation, room division and reconfiguration, decorated surfaces, paintings, and image-bearing surface use, while the weaker separation diagnostics and the smaller second cluster do not isolate perceptual extension as a distinct recovered contextual mode. Within the broad expository corpus used here, the analysis therefore supports fusuma as both a movable partition and an image-bearing architectural surface, consistent with accounts of fusuma as a mediator of domestic spatial relations [25]. The further claim that fusuma-e perceptually extends an otherwise confined interior remains an architectural interpretation informed by the pictorial and architectural literature [24,50].

4.8. Implications for Architectural Semantic Analysis

In this study, "recurring contextual patterns" refers to differentiated contextual structure recovered within individual terms, particularly where it remains identifiable in both FULL and CLEAN. Such recurrence does not, by itself, establish a common type of semantic ambiguity across the eight terms.

The proposed architectural readings remain term-specific: boundary and connection in engawa; symbolic and representational interiority in tokonoma; scenic incorporation in shakkei; relational breadth in ma; conceptual abstraction in mu; an interiorized threshold in doma; luminous permeability and flexible enclosure in shōji; and partitioning and spatial recomposition in fusuma. The shared register framework makes these readings comparable, while the strength of each interpretation depends on its corpus evidence and relevant architectural scholarship.

The computational contribution is to make these distinctions explicit and testable across paired corpus conditions. The computational analysis supplements architectural judgment and avoids establishing a final ontology of the terms. It shows where an architectural interpretation has recurring contextual support, where that support weakens under CLEAN conditions, and where the interpretation requires more caution than the theory alone might suggest. The FULL and CLEAN comparison is central because it examines whether contextual patterns persist after definitional support is reduced. It helps identify patterns that persist or weaken under CLEAN; interpreting those patterns as semantic ambiguity still requires representative contexts and architectural judgment.

4.9. Limitations and Scope

Several limitations arise from the study's design. The first is the narrow target set. The paper examines eight selected terms, not Japanese architectural vocabulary as a whole. This narrowness is methodologically useful because it keeps the analysis focused on local ambiguity, but it also limits generalizability. The findings provide a focused ambiguity analysis within the audited corpus, with claims limited to the selected terms and corpus conditions. Related work within the same research program has applied the broader computational-semantic framework to contextual comparison among selected architectural terms [43], indicating that the analytical approach can be extended beyond the within-term ambiguity problem examined here. Such extension, however, does not by itself establish generalizability to other vocabularies, languages, or architectural traditions, which would require separate empirical testing. It nevertheless highlights the potential of computational analysis to generate new insights into how architectural terms are interpreted, related, and compared. With further validation and extension, such approaches may inform broader investigations of architectural meaning and contribute to design-related inquiry by clarifying how spatial and conceptual categories are understood and articulated.

A second limitation concerns the English-language corpus. The findings describe uses of Japanese-origin terms in mixed English-language explanatory discourse, including some architectural writing, rather than Japanese-language usage or a single specialist architectural genre. This source mix fits the question about English-language explanation, but the results cannot establish Japanese-language meanings, translation equivalence, or patterns specific to specialist practice. A more discipline-specific corpus might sharpen some distinctions, especially for ma, shōji, and fusuma, or reveal different forms of overlap.

A further limitation concerns cluster extraction and granularity. Matched contexts refer to usable sentence-level target-span vectors, not raw mentions. This stricter procedure improves auditability, but it also means that the reported counts are not raw frequency totals. The active run also used a conservative binary ambiguity probe, so $k = 2$ indicates an accepted two-way contextual contrast, not a complete estimate of dictionary senses.

Future work could examine whether larger specialist corpora, alternative clustering criteria, or carefully audited fallback strategies produce finer-grained distinctions for terms such as ma, shōji, and fusuma. The 1.00 values in the screening audit are controlled audit outcomes within this pipeline, not proof of perfect semantic classification.

Finally, only engawa received a bespoke secondary concept-alignment test, so this asymmetry should remain explicit. We selected engawa and the boundary-edge, connector-flow, and riverbank probes from a prior architectural hypothesis, making this one-term analysis exploratory and theory-guided. We interpreted the other terms through the shared ambiguity pipeline. Accordingly, the additional concept-alignment evidence available for engawa is not directly equivalent to the evidential basis available for the other seven terms. Because the other terms did not receive comparable probes, the additional alignment results cannot rank engawa as intrinsically more ambiguous or semantically stronger than the other terms. Similar tests could be developed for tokonoma or shakkei, but only if the architectural hypothesis is precise enough to avoid becoming a loose confirmation exercise. These limits show that computational evidence strengthens architectural interpretation only when the interpretation remains constrained by the corpus and method.

5. Conclusions

This study examined whether a small, screened set of Japanese architectural terms exhibits recurring contextual-semantic structures after overt definitional support is reduced. Several terms retained interpretable contextual-semantic structures under CLEAN.

RQ1 asked whether selected Japanese architectural terms exhibit recurring contextual-semantic structures that support interpretation beyond singular material reference. The results provide a qualified affirmative answer. Several terms showed patterned contextual variation, indicating that their architectural significance cannot be reduced to stable object labels.

RQ2 concerned the durability of the identified structures across FULL and CLEAN. In most accepted cases, contextual-semantic structure persisted across both conditions, but its form and evidential strength varied. Engawa and tokonoma retained structure beyond simple material description, each with a one-corpus separation caveat; shakkei remained oriented toward represented/imagined alignment and perceptual incorporation; ma retained broad relational scope; and shōji retained stable constructional and permeability-related structure. Mu remained chiefly conceptual and lexical. Fusuma supported partitioning and spatial recomposition more securely than representational extension, although its separation diagnostics were weak. Doma remained a low-stability case despite complete target-span identification and 20 matched contexts in each condition. RQ2 therefore receives a qualified affirmative answer: persistence across corpus conditions does not imply uniform evidential strength or one common ambiguity structure.

RQ3 addressed whether the identified patterns could be interpreted through the five architectural registers: physical–spatial, symbolic–theoretical, represented/imagined, boundary/transition, and hybrid. The results can be interpreted with the five registers, provided their alignments are understood as conditional on theory-guided prompt choices rather than as categories discovered independently from the corpus. Their value lies in organizing comparison across terms, not in fixing the final meaning of those terms.

Lastly, RQ4 focused on engawa. The focused concept-alignment test identified positive alignment with both boundary-edge and connector-flow logic. Boundary-edge and riverbank alignment sat near the top of the raw concept-alignment plot, while connector-flow and derived hybridity also showed positive alignment. Taken together, the preselected probes yield exploratory alignment compatible with a combined boundary-and-connection reading of engawa; they do not establish that reading as a distinct semantic category. Building on this evi-

dence and the architectural literature, the paper proposes the modal littoral-circuit space formulation as an architectural interpretation of engawa: a shifting edge condition that connects interior and exterior by maintaining, thickening, and inhabiting the boundary between them.

The main conclusion is that several selected Japanese architectural terms show recurring contextual differentiation in the paired English-language corpus conditions, while the architectural interpretation and evidential strength of those patterns differ by term. Contextual computational analysis can make this variation more explicit and testable while remaining anchored in architectural interpretation. By distinguishing concise material definitions from recurring symbolic, relational, perceptual, and experiential uses, the framework may also inform layered architectural glossary and encyclopedia entries, provided that its outputs remain subject to architectural, historical, and linguistic review.

Author Contributions: Conceptualization, G.G. and S.Y.; methodology, G.G.; software, G.G.; validation, G.G. and S.Y.; formal analysis, G.G.; investigation, G.G.; resources, G.G.; data curation, G.G.; writing—original draft preparation, G.G.; writing—review and editing, G.G.; visualization, G.G.; supervision, S.Y.; project administration, G.G.; funding acquisition, S.Y. All authors have read and agreed to the published version of the manuscript.

Funding: The APC was funded by JSPS KAKENHI Grant Number JP23K04201 ("Co-Creation Between Deep Learning-Based Generative AI and Humans") and the Institute of Disaster Mitigation for Urban Cultural Heritage, Ritsumeikan University. The funders had no role in the design of the study, the collection, analysis, or interpretation of data, the writing of the manuscript, or the decision to publish the results.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All supporting materials for this study, including the scripts, configuration files, runtime log, generated CSV outputs, provenance files, manuscript figures, supplementary figures, and archived results from the reported run, are publicly available in GitHub release v1.0.0, released on 8 June 2026: https://github.com/ArchiScrub/Semantic-Ambiguity-in-Japanese-Architectural-Terms-A-Computational-Contextual-Analysis/releases/tag/v1.0.0 (accessed on 8 June 2026). These archived materials document the reported analysis and support the findings of this study.

Acknowledgments: During the preparation of this work, the author used ChatGPT 5.5 to support code writing, study-related reasoning, debugging, and code refinement; to assess and refine the interpretation of results to avoid overstatement; and to improve English clarity and readability. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Conflicts of Interest: The authors declare no competing interests.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
APC	Article processing charge
ARI	Adjusted Rand index
BERT	Bidirectional Encoder Representations from Transformers
BIM	Building information modeling
CC BY	Creative Commons Attribution
CLEAN	Corpus condition with glossary-like and explicitly definitional sentences removed
CSV	Comma-separated values
cuDNN	CUDA Deep Neural Network library
ELMo	Embeddings from Language Models

FULL	Corpus condition preserving broader expository, glossary-like, and definitional context
GPT-2	Generative Pre-trained Transformer 2
JSPS	Japan Society for the Promotion of Science
KAKENHI	Grants-in-Aid for Scientific Research
NLP	Natural language processing
RQ	Research question
TF32	TensorFloat-32
Word2Vec	Word to Vector
YAML	YAML Ain't Markup Language

Appendix A

Appendix A lists the eight target terms and provides brief dictionary-based definitions as baseline lexical references for the terms analyzed in the paper.

1. Doma (土間): Dirt floor; pit; parterre [51].
2. Engawa (縁側): (n) veranda; porch; balcony; open corridor [51].
3. Fusuma (襖): (n) fusuma (Japanese sliding screen) [51].
4. Ma (間): (n-adv, n) space (between); gap; interval; distance; (n-adv, n-t) time (between); pause; break; (n-adv, n, n-t) span (temporal or spatial); stretch; period (while); (n-adv, n) relationship (between, among); (n) space; room; time; pause; (n, suf) interval; period of time; among; between; inter- [51].
5. Mu (無): (n) nothing; naught; nought; nil; zero [51].
6. Shakkei (借景): (n) borrowed scenery (incorporation of a landscape element outside a garden in its design) [51].
7. Shōji (障子): (n) shōji (paper sliding door) [51].
8. Tokonoma (床の間): (n) tokonoma (alcove where art or flowers are displayed) [51].

References

1. Forty, A. *Words and Buildings: A Vocabulary of Modern Architecture*; Thames & Hudson: New York, NY, USA, 2004.
2. Pérez-Gómez, A. *Architecture and the Crisis of Modern Science*; MIT Press: Cambridge, MA, USA, 1983.
3. Harries, K. *The Ethical Function of Architecture*; MIT Press: Cambridge, MA, USA, 1998.
4. Pallasmaa, J. *The Eyes of the Skin: Architecture and the Senses*, 4th ed.; John Wiley & Sons: Hoboken, NJ, USA, 2024.
5. Hillier, B.; Hanson, J. *The Social Logic of Space*; Cambridge University Press: Cambridge, UK, 1984.
6. Amit, R. On the Structure of Contemporary Japanese Aesthetics. *Philos. East West* **2012**, *62*, 174–185. [CrossRef]
7. Ashihara, Y. *Exterior Design in Architecture*; Van Nostrand Reinhold: New York, NY, USA, 1970.
8. Inoue, M. *Space in Japanese Architecture*; Weatherhill: New York, NY, USA, 1985.
9. Sand, J. *House and Home in Modern Japan: Architecture, Domestic Space, and Bourgeois Culture, 1880–1930*, 1st ed.; Harvard University Asia Center: Cambridge, MA, USA, 2003; Volume 223.
10. Pilgrim, R.B. Intervals ("Ma") in Space and Time: Foundations for a Religio-Aesthetic Paradigm in Japan. *Hist. Relig.* **1986**, *25*, 255–277. [CrossRef] [PubMed]
11. Nitschke, G. MA: Place, Space, Void. *Kyoto J.* **2018**. Available online: https://kyotojournal.org/culture-arts/ma-place-space-void/ (accessed on 11 October 2025).
12. Isozaki, A. *Japan-Ness in Architecture*; Stewart, D.B., Ed.; MIT Press: Cambridge, MA, USA, 2006.
13. Marra, M. Japanese Aesthetics: The Construction of Meaning. *Philos. East West* **1995**, *45*, 367–386. [CrossRef] [PubMed]
14. Saito, Y. The Moral Dimension of Japanese Aesthetics. *J. Aesthet. Art Crit.* **2007**, *65*, 85–97. [CrossRef]
15. Fowler, M. Unfolding Architecture, Enfolding Landscape: The Shakkei at Geppa-Rō Pavilion. *Archit. Theory Rev.* **2023**, *27*, 235–257. [CrossRef]
16. Kuitert, W. Borrowing Scenery and the Landscape That Lends—The Final Chapter of Yuanye. *J. Landsc. Archit.* **2015**, *10*, 32–43. [CrossRef]
17. Kurokawa, K. Each One A Hero: The Philosophy of Symbiosis, Chapter 8: Intermediary Space. Available online: https://www.kisho.co.jp/page/305.html (accessed on 12 May 2026).
18. Wong, J.F. Conceptual Engawa: The Experience of Ring-Based Circulation in Tadao Ando Museums. *Environ. Plan. B Plan. Des.* **2014**, *41*, 229–250. [CrossRef]

19. Hou, Z.; Kong, S.; Wang, Y.; An, Q. The Engawa as Spatial Mediator: Transformation of Design Mechanisms in Japanese Teahouses. *Buildings* **2026**, *16*, 1113. [CrossRef]
20. Yang, Y. *Shikii* in Defining Japanese "Public" Space: Case Study of Shinonome Canal Court in Tokyo. *J. Urban Manag.* **2026**, *15*, 421–435. [CrossRef]
21. Anderson, J.L. Japanese Tea Ritual: Religion in Practice. *Man* **1987**, *22*, 475–498. [CrossRef] [PubMed]
22. Kondo, D. The Way of Tea: A Symbolic Analysis. *Man* **1985**, *20*, 287–306. [CrossRef] [PubMed]
23. Kumakura, I. *Japanese Tea Culture: The Heart and Form of Chanoyu*; Japan Publishing Industry Foundation for Culture: Tokyo, Japan, 2023.
24. Stoichita, V.I. *The Self-Aware Image: An Insight into Early Modern Meta-Painting*; Cambridge University Press: Cambridge, UK, 1997.
25. Ingarden, K.; Watanabe, Y.K. The Japanese Traditional House: Spatial Representation of Dualistic Social Codes. *Hous. Environ.* **2025**, *54*, 25–37. [CrossRef]
26. Kim, J. A Reexamination of Sekkyaku-Honi (接客本位) in Edo Middle-Lower Samurai Residences from a Space Syntax Perspective: Reinterpreting Spatial Evolution in Modern Japanese Urban Housing. *J. Asian Archit. Build. Eng.* **2026**, *25*, 4097–4129. [CrossRef]
27. Nuijsink, C. From '"Container"' to '"Lifestyle:"' Kazuyo Sejima, Sou Fujimoto and the Destruction of the Nuclear Family Box. *Interiors* **2021**, *11*, 132–156. [CrossRef]
28. Ding, Y.; Ma, J.; Luo, X. Applications of Natural Language Processing in Construction. *Autom. Constr.* **2022**, *136*, 104169. [CrossRef]
29. Locatelli, M.; Seghezzi, E.; Pellegrini, L.; Tagliabue, L.C.; Di Giuda, G.M. Exploring Natural Language Processing in Construction and Integration with Building Information Modeling: A Scientometric Analysis. *Buildings* **2021**, *11*, 583. [CrossRef]
30. Horvath, A.-S. How We Talk(Ed) about It: Ways of Speaking about Computational Architecture. *Int. J. Archit. Comput.* **2022**, *20*, 150–175. [CrossRef]
31. Ono, K.; Tanikawa, N.; Yamada, S. The Image of a Famous Architect as a Whole as It Appears in Academic Discourse: Comparison of Conceptual Diversity and Interpretation Using Text Mining. In *Proceedings of the Architectural Informatics: Proceedings of the 30th International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA 2025), Volume 1*; Association for Computer-Aided Architectural Design Research in Asia (CAADRIA): Tokyo, Japan, 2025; Volume 1, pp. 561–570.
32. Suissa, O.; Elmalech, A.; Zhitomirsky-Geffet, M. Text Analysis Using Deep Neural Networks in Digital Humanities and Information Science. *J. Assoc. Inf. Sci. Technol.* **2022**, *73*, 268–287. [CrossRef]
33. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention Is All You Need. In *Proceedings of the Advances in Neural Information Processing Systems 30 (NIPS 2017), Long Beach, CA, USA, 4–9 December 2017*; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Vol. 30, pp. 5998–6008.
34. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*; Burstein, J., Doran, C., Solorio, T., Eds.; Association for Computational Linguistics: MN, USA, 2019; pp. 4171–4186.
35. Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient estimation of word representations in vector space. *arXiv* **2013**, arXiv:1301.3781.
36. Erk, K. Vector Space Models of Word Meaning and Phrase Meaning: A Survey. *Lang. Linguist. Compass* **2012**, *6*, 635–653. [CrossRef]
37. Navigli, R. Word Sense Disambiguation: A Survey. *ACM Comput. Surv. (CSUR)* **2009**, *41*, 1459355. [CrossRef]
38. Ethayarajh, K. How Contextual Are Contextualized Word Representations? Comparing the Geometry of BERT, ELMo, and GPT-2 Embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*; Inui, K., Jiang, J., Ng, V., Wan, X., Eds.; Association for Computational Linguistics: Hong Kong, China, 2019; pp. 55–65.
39. Rogers, A.; Kovaleva, O.; Rumshisky, A. A Primer in BERTology: What We Know About How BERT Works. *Trans. Assoc. Comput. Linguist.* **2020**, *8*, 842–866. [CrossRef]
40. Pineau, J.; Vincent-Lamarre, P.; Sinha, K.; Lariviere, V.; Beygelzimer, A.; d'Alche-Buc, F.; Fox, E.; Larochelle, H. Improving Reproducibility in Machine Learning Research (A Report from the NeurIPS 2019 Reproducibility Program). *J. Mach. Learn. Res.* **2021**, *22*, 1–20.
41. Gjata, G.; Yamada, S. Reproducibility and Validation of a Computational Framework for Architectural Semantics: A Methodological Study with Japanese Architectural Concepts. *Buildings* **2025**, *15*, 4107. [CrossRef]
42. Gjata, G.; Yamada, S. Quantitative Mapping of Conceptual Hierarchies and Data-Driven Taxonomies of Japanese Architectural Concepts: A 28-Term Testbed. *Architecture* **2026**, *6*, 62. [CrossRef]
43. Gjata, G.; Yamada, S. Semantic Comparison of Architectural Terms: A Computational Contextual Analysis. *Architecture* **2026**, *6*, 121. [CrossRef]
44. Morse, E.S. *Japanese Homes and Their Surroundings*; Ticknor and Company: Boston, MA, USA, 1886.

45. Parent, M.N. Amado 雨戸. In *JAANUS: Japanese Architecture and Art Net Users System*; JAANUS: Tokyo, Japan, 2001; Available online: https://www.aisf.or.jp/ jaanus/deta/a/amado.htm (accessed on 25 March 2026).
46. Hojda, O. Japan and Modern Architecture, 1945–1970. Discourse in the Mid-20th-Century Europe. 1960s: Emancipation, Criticism. *Auc. Philos. Hist.* **2021**, *2020*, 195–205. [CrossRef]
47. The Metropolitan Museum of Art Brush Writing in the Arts of Japan. Available online: https://www.metmuseum.org/exhibitions/listings/2013/brush-writing (accessed on 15 September 2025).
48. Pilgrim, R.B. The Artistic Way and the Religio-Aesthetic Tradition in Japan. *Philos. East West* **1977**, *27*, 285–305. [CrossRef] [PubMed]
49. Jiang, P. The Ambiguous Boundary between Buildings and Environment: Exploring the Interpretation of Traditional Aesthetic in Contemporary Japanese Architecture. *Open J. Civ. Eng.* **2022**, *12*, 415–437. [CrossRef]
50. Yamamoto, H. About the Fusuma Paintings from Obai-in Temple. Available online: https://www.kyohaku.go.jp/eng/learn/home/dictio/kaiga/oubai/ (accessed on 26 May 2026).
51. Hai, B.H. RomajiDesu English Japanese Dictionary. Available online: http://www.romajidesu.com/dictionary/ (accessed on 7 November 2025).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.