

Article

Evidence-Aware Data Model for Explainable Post-Hurricane Building Damage Assessment

Ahram Song ¹, Jinha Jung ², Anjin Chang ³ and Seula Park ^{4,*}

¹ Department of Location-Based Information System, Kyungpook National University, Sangju-si 37224, Gyeongsangbuk-do, Republic of Korea; ars@knu.ac.kr

² School of Civil and Construction Engineering, Purdue University, West Lafayette, IN 47907, USA; jinha@purdue.edu

³ Department of Biosystems and Agricultural Engineering, College of Agriculture and Natural Resources, Michigan State University, East Lansing, MI 48824, USA; changanj@msu.edu

⁴ Department of Civil and Environmental Engineering, Gachon University, Seongnam-si 13120, Gyeonggi-do, Republic of Korea

* Correspondence: seulapark@gachon.ac.kr

Abstract

Recent learning-based damage assessment methods using remote sensing data have greatly contributed to the rapid identification and classification of damaged buildings from pre- and post-disaster imagery. However, most outputs are provided in the form of final damage labels or scores. To address this limitation, this study proposes an evidence-aware data model for explainable post-disaster building damage assessment. The proposed model is derived from a component-level roof-damage assessment workflow using optical imagery and post-disaster surface-height data, and is designed to describe the observational evidence, intermediate indices, and decision rules leading to a specific damage grade, as well as its interpretation as a post-disaster change event. Based on the proposed logical data model, a hybrid spatial-graph database was implemented to support both spatial retrieval and the semantic explanation of damage assessment results. The proposed model was evaluated through six competency questions derived from different damage assessment requirements. The results confirmed that the proposed model can support component-level retrieval of damage results, tracing of damage evidence and assessment rules, spatial visualization of damaged building components, and integrated spatial-semantic explanation for individual objects. The results provide a foundation for managing disaster damage assessment processes and results as explainable and reusable structured data.

Keywords: evidence-aware data model; data modeling; post-disaster damage assessment; hybrid database



Academic Editor: Wolfgang Kainz

Received: 3 July 2026

Revised: 6 September 2026

Accepted: 7 September 2026

Published: 10 September 2026

Copyright: © 2026 by the authors.

Published by MDPI on behalf of the

International Society for

Photogrammetry and Remote Sensing.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Post-disaster building damage assessment is a critical task for emergency response, recovery planning, field inspection prioritization, and spatial database updating. Remote sensing has long been used to support this task because satellite, aerial, unmanned aircraft systems (UAS), synthetic aperture radar (SAR), light detection and ranging (LiDAR), and elevation-derived data can provide rapid and spatially extensive observations of affected areas [1,2]. Earlier studies demonstrated the value of remote sensing for identifying collapsed or damaged buildings after earthquakes and other disasters, while also noting limitations related to data availability, viewing geometry, damage complexity, and interpretation uncertainty [1,2].

Recent developments have shifted the field from manual interpretation and object-based analysis toward large-scale benchmark datasets and learning-based damage classification. For example, xBD provided pre- and post-disaster satellite imagery with building polygons and ordinal damage labels, enabling systematic training and the evaluation of damage assessment models across multiple disaster types [3]. SpaceNet 8 further expanded disaster-related remote sensing benchmarks by combining building, road, and flood-related tasks [4]. More recent multimodal datasets have introduced optical and SAR imagery to support damage assessment under more diverse observation conditions [5]. Along with these datasets, convolutional neural networks, Siamese architectures, transformers, and multi-view fusion models have improved the localization and classification of damaged buildings from pre- and post-disaster imagery [6,7]. These studies have improved the accuracy and scalability of damage mapping, but their outputs are commonly represented as final labels, class probabilities, or damage scores.

While image-based learning methods provide powerful tools for rapid damage mapping, roof and structural damage are not always fully characterized by optical appearance changes alone. Damage may involve height loss, roof-surface flattening, structural deformation, debris accumulation, and the fragmentation of surface elements, which require additional geometric or surface-related information for interpretation. UAS-based studies have used very high-resolution oblique imagery, point clouds, object-based image analysis, and semantic reasoning to capture detailed structural-damage patterns [8,9]. LiDAR- and digital surface model (DSM)-based studies have used height, roof-area, volume, orientation, and shape-related indicators to quantify building and roof damage after disasters [10]. These studies show that structural surface information can complement optical cues. However, the indicators, thresholds, and reasoning procedures used in such approaches are often embedded in analysis scripts, raster products, vector attributes, or methodological descriptions, rather than being formalized as reusable data objects linked to damage assessment results.

Considering the need to support the sustained use of damage assessment processes and outputs, it is increasingly important that damage assessment results derived from multimodal datasets be managed not merely as final outputs, but as structured data models and operational databases. This approach enables the integrated management of various observational data and analytical results. In the disaster management domain, comprehensive databases are required to accommodate disaster-specific information and to satisfy the information requirements of different phases of disaster management [11]. In this context, beyond damage-detection algorithms, previous studies have also attempted to use remote sensing-derived damage information for geospatial database updating or to structure disaster-related damage information using geospatial data models. Post-disaster building database updating has been investigated by integrating pre-disaster building records, such as OpenStreetMap data, with multi-temporal satellite observations [12]. This research demonstrates the importance of linking damage assessment outputs to spatial database maintenance, but it primarily focuses on identifying changed or damaged building records for updating, rather than representing the evidence, index, rule, and change-inference logic underlying each update decision. In flood damage assessment, BIM-GIS integration has been proposed to support the micro-level assessment and 3D visualization of building damage by representing flood parameters, building components, materials, utilities, and damage costs within a unified data model [13,14].

Existing standards for the semantic and geometric representation of spatial objects can also support the systematic structuring of post-disaster damage assessment processes and results. For example, CityGML application domain extensions (ADEs) have also been used to extend generic 3D city models with domain-specific attributes and semantic in-

formation [15,16]. Based on this mechanism, CityGML-based approaches have also been developed to integrate multi-source urban flooding data and simulation results using application domain extensions, thereby improving the data interoperability, semantic representation, and 3D visualization for flood-related analysis [17,18]. In post-earthquake applications, UAS-derived point clouds have been used to generate LOD3 CityGML building models and enrich individual building components, such as walls and roofs, with seismic damage semantics [19]. These studies demonstrate that semantic geospatial data models can improve the representation, visualization, and management of disaster-related damage information. However, existing approaches remain focused mainly on database updating, component representation, simulation-data integration, 3D visualization, or damage grade attribution. In other words, they primarily represent updated records, simulation objects, or resulting damage states, rather than the process-oriented assessment workflow through which those states are derived. Additionally, GeoSPARQL supports the semantic representation and querying of spatial objects and their relationships [20], making it useful for structuring the spatial context among urban objects in disaster scenarios. Furthermore, existing provenance and observation standards provide relevant concepts for structuring post-disaster damage assessment processes. PROV-O, a provenance ontology, provides general concepts for tracing data provenance and derivation. It can be specialized by defining new classes and properties to model provenance information for different applications and domains [21]. The Semantic Sensor Network (SSN) ontology provides a conceptual framework for modeling observations, procedures, and features of interest. It is designed to support a wide range of applications and use cases, including satellite imagery and large-scale scientific monitoring [22]. However, these standards are primarily designed for general-purpose representations of provenance, observations, and spatial objects, rather than for explicitly modeling domain-specific assessment decision processes. Although their concepts can be extended or specialized for different domains, they do not prescribe the sequence through which evidence derived from multi-source observations is integrated, evaluated against assessment rules and conditions, and translated into final damage assessment results. Accordingly, a more assessment-specific representation is required to explicitly capture the decision process and underlying rationale associated with post-disaster damage assessment outcomes.

More broadly, recent explainable AI and knowledge graph studies have emphasized the importance of interpretable decisions and relation-based information representation [23–25]. However, these general frameworks do not directly define the assessment-specific entities and relationships required to explain how remote sensing observations are converted into evidence records, summarized into continuous damage indices or indicators, interpreted through assessment rules, and connected to damage results and inferred change events. This limitation is particularly relevant to assessment traceability: a final label or score alone does not preserve the supporting evidence, data-availability conditions, or rule conditions associated with the result. In existing remote sensing-based damage assessment workflows, such information is often distributed across raster products, vector attributes, tabular outputs, thresholds, and analysis scripts. Therefore, an explicit evidence-aware data model is still needed to connect observation-derived evidence, damage indices, assessment rules, damage results, and change events within a single explainable and queryable workflow. Such data models are essential for ensuring the interpretability of damage assessment results by incorporating semantic relationships among assessment information. They also support inference by preserving structured assessment rules and intermediate information derived during the assessment process. Furthermore, they provide a basis for operating a standardized damage assessment framework by managing assessment results within a structured database, rather than simply storing them as numerical values.

To address this gap, this study proposes an evidence-aware data model for explainable post-disaster building damage assessment. Explainability in this study does not refer to explaining the internal mechanisms of the learned model. Instead, it refers to the traceability of a damage assessment result through its supporting evidence and applied assessment rule. The model is derived from a component-level roof-damage evidence assessment workflow using optical imagery and post-disaster surface-height data as observational sources. In the workflow, surface and optical evidence records are summarized into a continuous roof-damage evidence index, which can then be interpreted through assessment rules to produce damage results and inferred change events. The focus of this paper is not on proposing a new damage index, but on formalizing the data objects and relationships required to store, query, explain, and operationalize an evidence-based damage assessment workflow.

The proposed model addresses three requirements of post-disaster damage assessment: explaining how damage results are derived, inferring change events from rule-based results, and managing assessment outputs in an appropriate database structure. Although the implementation in this paper focuses on hurricane-induced roof damage, the model is designed to be extensible to other disaster types, evidence sources, damage indices, and assessment rules at the schema level, rather than implying empirical validation across multiple disaster types. To support these requirements, the model is implemented using a hybrid spatial-graph database structure. Spatial geometries and component-level assessment units are managed in a spatial database, whereas evidential, rule-based, and inferential relationships are represented in a property graph database to support multi-hop relation tracing. This structure enables spatial querying, object-level visualization, and the relation-based explanation of the assessment process, thereby supporting traceable verification and more reliable decision-making in post-disaster damage assessment.

The main contributions of this study are as follows. First, it formalizes an evidence-aware data model that connects observation-derived evidence, damage indices, assessment rules, damage results, and inferred change events. Rather than proposing a new general-purpose provenance standard, the proposed model explicitly structures an evidence- and rule-based assessment lineage for post-disaster damage assessment. Second, it demonstrates a hybrid spatial-graph database implementation that combines spatial object management with multi-hop explanation tracing. Third, it evaluates the implemented model through competency questions covering damage-result retrieval, evidence tracing, rule reconstruction, severe-damage mapping, and hybrid spatial-graph querying.

2. Data Model for Explainable Disaster Assessment

2.1. Roof-Damage Evidence Assessment Workflow

The proposed data model was derived from a component-level roof-damage assessment workflow used to populate the case-study database. This workflow represents one case-specific instantiation of the model rather than a universally applicable damage assessment method. An assessment unit is defined as

$$u_i = (b_i, c_i) \quad (1)$$

where b_i denotes a building object and c_i denotes the assessed component, which corresponds to the roof in this case study. The pre-disaster building footprint provides the spatial reference, and an inner evidence region is used to reduce boundary contamination and residual misregistration.

Two complementary evidence streams are used: surface evidence derived from post-disaster surface-height data and optical evidence derived from pre- and post-disaster imagery. Evidence quantities stored in the data model are represented on a normalized

[0, 1] scale, where larger values indicate stronger evidence consistent with roof damage. For descriptors requiring scaling, robust normalization is defined as

$$\tilde{x} = \text{clip}\left(\frac{x - Q_5}{Q_{95} - Q_5}, 0, 1\right), \quad (2)$$

where Q_5 and Q_{95} denote the fifth and 95th percentiles of the reference distribution. These values are used only as robust scaling bounds and not as damage-classification thresholds. Corresponding pre- and post-disaster optical descriptors are normalized using a common pooled reference distribution.

2.1.1. Surface Evidence

Surface evidence characterizes the post-disaster geometric condition of the roof. Height-related evidence H_i combines absolute height-distribution evidence H_i^{abs} , reflecting loss of elevated surfaces and low-height dominance, with within-object height-depression evidence H_i^{dep} :

$$H_i = 1 - (1 - H_i^{abs})(1 - H_i^{dep}). \quad (3)$$

Surface-pattern evidence incorporates residual surface roughness and the fragmentation of elevated regions. Let ρ_i^{res} denote residual roughness after the removal of the dominant planar surface and $\mathcal{N}(\rho_i^{res})$ its normalized value. Fragmentation is represented as $\kappa_i = 1 - (A_i^{max} / A_i^{high})$, where A_i^{high} is the total elevated-surface area and A_i^{max} is the largest connected elevated-surface component; $\kappa_i = 0$ when no elevated component is present. Surface-pattern evidence is then defined as

$$G_i = \max[H_i \mathcal{N}(\rho_i^{res}), \kappa_i] \quad (4)$$

Finally, the height-related and surface-pattern evidence are combined using evidence union:

$$S_i = 1 - (1 - H_i)(1 - G_i). \quad (5)$$

The resulting $S_i \in [0, 1]$ represents composite surface evidence of roof-surface disruption. Accordingly, the surface-evidence stream captures complementary information related to elevated-surface loss, within-object height depression, residual roughness, and fragmentation. While the primitive height and surface descriptors are intermediate quantities used in evidence generation, H_i , G_i , and S_i are retained as explicit DamageEvidence records in the proposed data model.

2.1.2. Optical Evidence

Optical evidence is calculated at the assessment-unit level rather than through direct pixel-wise subtraction because the pre- and post-disaster observations differ in acquisition and imaging characteristics. Three object-level descriptors are used, namely color variability $C_{i,t}$, texture complexity $T_{i,t}$, and edge strength $B_{i,t}$, where t denotes the acquisition time. Corresponding descriptors are normalized using the common pre/post reference distributions described above. Post-disaster optical disruption and temporal optical-change evidence are defined as

$$O_i^{post} = \frac{\tilde{C}_{i,post} + \tilde{T}_{i,post} + \tilde{B}_{i,post}}{3}, \quad (6)$$

$$\Delta O_i = \frac{|\tilde{C}_{i,post} - \tilde{C}_{i,pre}| + |\tilde{T}_{i,post} - \tilde{T}_{i,pre}| + |\tilde{B}_{i,post} - \tilde{B}_{i,pre}|}{3} \quad (7)$$

The final optical evidence combines post-disaster disruption and temporal change as

$$O_i = O_i^{post} \Delta O_i, \quad (8)$$

This formulation requires both post-disaster visual disruption and corresponding temporal change to produce strong optical evidence. The quantities O_i^{post} , ΔO_i , and O_i are retained separately as DamageEvidence records.

2.1.3. Evidence Integration and Interpretation

Surface and optical evidence are integrated into the continuous roof-damage evidence index I_i . Let $U_i = \max(S_i, O_i)$, $M_i = (S_i + O_i)/2$, where U_i represents the stronger of the surface and optical evidence and M_i represents their mean. The continuous evidence index is then defined as

$$I_i = U_i - (U_i - M_i)(1 - U_i). \quad (9)$$

This integration retains strong evidence from either source while moderating the index when surface and optical evidence disagree. No learned weighting parameters are used. The resulting $I_i \in [0, 1]$ is interpreted as a continuous roof-damage evidence index rather than as an intrinsic categorical structural-damage grade.

For traceability, the evidence records are retained separately:

$$E_i = \{H_i, G_i, S_i, O_i^{post}, \Delta O_i, O_i\}. \quad (10)$$

The continuous index I_i is represented separately as a DamageIndex object. Table 1 summarizes the observation dependencies, conceptual meanings, immediate derivation dependencies, and database representations of the evidence and index quantities.

The surface and optical evidence streams remain separate after index generation because they may have different observational reliability and may also participate differently in the subsequent decision process. In the present case study, I_i provides the baseline categorical interpretation, while S_i can additionally participate in evidence-specific rule conditions.

When a categorical result is required, the stored evidence and index are evaluated using the applicable assessment-rule set:

$$D_i = f^{rule}(I_i, E_i; \mathcal{R}), \quad (11)$$

where D_i denotes the categorical damage grade and $\mathcal{R}$ denotes the applicable set of assessment rules. Each rule consists of explicit conditions and corresponding actions. The baseline categorical interpretation can be based on I_i , while individual evidence records in E_i may participate directly in rule conditions when required by the assessment workflow. The rule set also defines the precedence used when more than one condition is applicable.

DamageEvidence, DamageIndex, AssessmentRule, and DamageResult are represented as distinct but related information objects, allowing the applied evidence and decision rule to be traced independently. Evidence-specific conditions may modify the baseline index-based interpretation according to the rule configuration.

Missing observations are represented separately from low evidence values; zero denotes available but weak evidence rather than missing information. The present case-study configuration requires both surface and optical evidence and should not be applied unchanged when a required evidence source is unavailable. The resulting DamageResult may subsequently be associated with an InferredChangeEvent and an update readiness state, completing the traceable evidence–index–rule–result–change workflow.

Table 1. Evidence and index information represented in the reference roof-damage assessment workflow.

Symbol	Group	Observation Dependency	Conceptual Meaning	Derivation Dependency	Evidence Record Name
H_i	Surface	Post-disaster surface-height data	Height-related evidence associated with elevated roof-surface loss, low-height dominance, and within-object height depression	Surface-height observation	Height_Deficit_Hi
G_i			Surface-pattern evidence associated with surface roughness and fragmentation of elevated roof regions		Surface_Disorder_Gi
S_i			Composite geometric evidence describing roof-surface disruption	H_i, G_i	Surface_Composite_Si
O_i^{post}	Optical	Post-disaster optical imagery	Post-disaster visual disruption represented by object-level color, texture, and edge characteristics	Post-disaster optical descriptors	Optical_Disruption_Opost
ΔO_i		Pre- and post-disaster optical imagery	Temporal change between corresponding object-level optical characteristics	Pre/post optical descriptors	Optical_Change_DeltaO
O_i			Composite optical support for roof-damage interpretation	$O_i^{post}, \Delta O_i$	Optical_Support_Oi
I_i	Index	Surface and optical evidence	Continuous roof-damage evidence index	S_i, O_i	DamageIndex

2.2. The Proposed Data Model

In this study, a logical data model was designed to represent and store explainable disaster damage assessment information. The proposed model aims not only to represent the spatial distribution of damage and the spatial objects subject to assessment, but also to describe the evidential basis and logical process underlying each damage assessment result. In the proposed data model, a traceable rule-based explanation is considered complete when the assessment result can be linked to the damage index used in the decision, supporting evidence, the assessment inputs from which the evidence was derived, and the applied assessment rule, including its conditions and action. To this end, the proposed model was organized into four modules: the disaster event module, building object module, damage assessment module, and change-inference module. The disaster event module represents the disaster event that triggered damage assessment, while the building object module describes buildings and building components that serve as the targets of disaster damage assessment. The damage assessment module is the core module for supporting the traceability of assessment result, as it represents what was observed as damage evidence and how damage results were derived from such evidence. The change-inference module represents the inferred changes in building conditions after a disaster, based on the damage assessment results.

The detailed entities and attributes of each module were derived by considering three requirements: the post-disaster damage assessment workflow, the spatial query requirements for building-level and component-level management, and the explanation requirements for tracing how assessment results are derived. Based on these requirements, the relational-spatial data model was completed as summarized in Table 2. The DisasterEvent entity provides basic information on the disaster that caused the assessment to be performed, whereas Building and BuildingComponent provide information on spatially identifiable damage targets and assessment objects. The BuildingComponent entity is not restricted to roof components; when corresponding observations and evidence are available, different building components, such as the roof, facade, and ground/first-floor components, can be instantiated as separate assessment targets. In the Building entity, the *osm identifier* attribute was included to link with external building data, such as OpenStreetMap. AssessmentInput stores detailed information on the input data used for assessment, and DamageAssessment stores basic information on the damage assessment activity itself. In particular, the *assessment method*, *confidence*, and *evaluation indicator* attributes were defined to indicate whether the assessment was performed through automated detection or another assessment method, as well as the assessment-level reliability information and the specific indicator to which that confidence value applies.

To ensure the explainability of damage assessment, the proposed model represents the process of deriving damage results through DamageEvidence, DamageIndex, and AssessmentRule. As described in Section 2.1, the damage-related observations derived from input data, including H_i , G_i , S_i , O_i^{post} , ΔO_i , O_i , are stored as DamageEvidence. This entity stores the *evidence type*, *evidence name*, and *evidence value* attributes, which, respectively, describe the category of evidence, such as optical or surface, the specific evidence variable, and its observed or calculated value. Quantitative damage indicators computed from combinations of these evidence elements are stored in DamageIndex, and these indicators serve as intermediate decision values for deriving assessment grades. The *metric code* attribute in DamageIndex indicates the specific type of damage indicator used as the damage index, such as I_i , and enables different damage indicators to be distinguished within the assessment process.

Table 2. Entities and their properties in the relational-spatial data model.

Module Category	Entity	Property
Disaster event module	DisasterEvent	Event identifier, event name, disaster type, start date, end date, region, description
Building object module	Building	Building identifier, osm identifier, geometry
	BuildingComponent	Component identifier, component type, building identifier, geometry
Damage assessment module	AssessmentInput	Input identifier, input type, data source, acquisition date, pre/post flag
	DamageAssessment	Assessment identifier, event identifier, assessment date, assessment method, confidence, evaluation indicator
	DamageIndex	Index identifier, assessment identifier, metric code, metric name, index value
	DamageResult	Result identifier, assessment identifier, building identifier, component identifier, damage type, damage grade, field reference availability, valid from, valid to
	AssessmentRule	Rule identifier, rule version, rule name, rule type, rule priority, valid from, valid to, logic operator
	RuleCondition	Condition identifier, rule identifier, rule version, metric code, lower value, upper value, lower inclusive, upper inclusive
	RuleAction	Action identifier, rule identifier, rule version, damage grade code, change type code, update readiness code
	DamageEvidence	Evidence identifier, assessment identifier, evidence type, evidence name, evidence value
Change inference module	InferredChangeEvent	Change identifier, result identifier, change type, update readiness, inferred at

The final damage grade is determined according to the damage indices and predefined conditions and thresholds. The logical conditions used to derive the result are represented through the combination of AssessmentRule, RuleCondition, and RuleAction. Specifically, AssessmentRule defines the purpose and logical structure of each assessment rule. RuleCondition represents the threshold conditions that determine whether a rule is satisfied. Its attributes, such as *metric code*, *lower and upper value*, *lower and upper inclusive*, specify which metric is evaluated and how the applicable value range is defined. RuleAction represents the action executed when the rule conditions are satisfied. The *damage grade code*, *change type code*, and *update readiness code* attributes define the resulting damage grade, the inferred change type, and the level of readiness for existing database updating. To explicitly manage the history of changes in assessment rules, the AssessmentRule entity includes the *rule version*, *valid from*, and *valid to* attributes. These attributes enable the identification of the rule version applicable at a specific assessment time and the reconstruction of the rules and conditions on which past assessment results were based.

DamageResult is an entity that stores the final damage grade. It also stores the identifiers of the assessed building and component, which enables the direct retrieval of damage results from the spatial database. In addition, the *damage type* and *damage grade* attributes describe the type and severity of the assessed damage, while *field reference availability* indicates whether an independent field reconnaissance record is available for the corresponding assessment unit. DamageResult includes the *valid from* and *valid to* attributes to distinguish the validity periods of individual results when repeated assessments are performed for the same building component. A current result, indicated by *valid_to* IS

NULL, refers to the *DamageResult* that represents the most recently recorded assessment state for a given *BuildingComponent*.

InferredChangeEvent represents a post-disaster change hypothesis inferred from the damage result. The *change type* attribute describes the rule-based change hypothesis associated with the result, such as *NoChange*, *MinorDamage*, *MajorDamage*, *SevereDamage*, or the component-specific *RoofCollapseCandidate*. The *update readiness* attribute represents the operational state of the inferred change within a downstream database-update workflow. In the present model, this state distinguishes cases for which no update is required, additional evidence is required, manual review is required, the predefined criteria for update processing are satisfied, or the assessment reliability is insufficient. Importantly, *ready_to_update* indicates eligibility for downstream update processing under the defined rule criteria and does not itself indicate independent object-level validation or automatic modification of an external database.


In summary, the assessment context represented by *DisasterEvent*, *DamageAssessment*, and *AssessmentInput* defines which disaster was assessed and under what conditions the assessment was performed. *Building* and *BuildingComponent* represent the damaged buildings and their detailed components, which constitute the assessment targets. *DamageEvidence* and *DamageIndex* define the intermediate evidence and indicators used to derive the assessment results, while *AssessmentRule*, *RuleCondition*, and *RuleAction* define the criteria by which the final damage grade is determined. Finally, *DamageResult* defines the final damage grade assigned to each building component through the assessment flow, and *InferredChangeEvent* defines building change information that links the assessment result to database updating.

To explicitly represent the logical process through which assessment results are derived and to ensure the explainability of the proposed model, this study structured the relationship information among major entities as a separately labeled property graph (LPG) model. Although a relational-spatial data model can retrieve related records by combining tables through joins, it has limitations in explicitly representing multi-hop semantic paths from assessment results to supporting evidence, damage indices, assessment rules, and inferred change events in a reusable form. Therefore, this study defined 13 labeled relationships, focusing on the semantic relationships that are repeatedly queried in the process of retrieving and interpreting assessment results, as well as the core relationships required to support evidence traceability and rule-based explanation, as summarized in Table 3. These relationships are not intended to cover all possible combinations among entities. Rather, they were designed as core relationships that support evidence tracing, rule-based explanation, and the connection between spatial assessment targets and their corresponding damage assessment results. An LPG database allows new nodes and relationships to be added without redesigning the entire existing schema, making it relatively easy to extend the relationship structure. This extensibility provides an additional advantage of the LPG model in the hybrid architecture, enabling the relationship structure to be expanded as needed while PostGIS supports spatial data processing.

Figure 1 presents the Unified Modeling Language (UML) class diagram of the hybrid data model, integrating the relational-spatial schema and labeled property graph representation. The formal schema definition is provided in the Supplementary Materials. Fundamentally, the relational-spatial data model manages entity information that requires object-level spatial querying and mapping, whereas the property graph model stores explainable relationships among the core entities involved in damage assessment. This configuration was designed to minimize redundancy while balancing spatial queryability and explainability.

Table 3. Labeled relationships in the property graph model.

Head	Relationship	Tail	Semantics
DamageEvidence	observed_for	BuildingComponent	Associates damage evidence with the building component on which it was observed.
InferredChangeEvent	inferred_by	DamageResult	Links an inferred change candidate to the damage result from which it was derived.
DamageResult	derived_from	AssessmentRule	Links a damage grade to the assessment rule used for its derivation.
DamageEvidence	generated_in	DamageAssessment	Associates damage evidence with the batch damage assessment in which it was generated.
DamageIndex	computed_from	DamageEvidence	Links a damage index to the damage evidence elements used for its computation.
DamageResult	decided_by	DamageIndex	Links a damage result to the damage index on which it is based.
DamageAssessment	conducted_using	AssessmentInput	Associates a damage assessment with the input data used for the assessment.
AssessmentRule	has_condition	RuleCondition	Links an assessment rule to the rule conditions that define its applicability.
AssessmentRule	triggers	RuleAction	Links rule to the rule action executed when the conditions are satisfied.
Building	composed_of	BuildingComponent	Represents the composition relationship between a building and its components.
DamageResult	result_for	BuildingComponent	Associates a damage result with the building component to which the result is assigned.
DamageResult	revises	DamageResult	Links a reassessment result to the previous damage result that it revises.
DamageEvidence	extracted_from	AssessmentInput	Links damage evidence to the assessment input(s) from which it was extracted.

**Figure 1.** UML class diagrams of the proposed hybrid data model.

Specifically, the relationship between *DamageEvidence* and *DamageIndex* is not managed with a foreign-key in the relational-spatial data model, but is instead represented by the *computed_from* relationship in the property graph model. This is because a single damage index can be derived from a combination of multiple pieces of evidence, forming a many-to-many relationship that is more appropriately represented as a graph-based explanatory relationship.

In the same context, the relationships between *DamageResult*, *DamageIndex*, and *AssessmentRule* are treated as inferential relationships that explain which indices and rules led to a given assessment result. Therefore, rather than representing these relationships through foreign keys in the relational-spatial data model, they are expressed in the property graph model using the *decided_by* and *derived_from* relationships. For assessment rules, the relational-spatial data model stores the rule definitions, whereas the property graph model represents the application relationships between damage results and rules. In addition, *DamageResult* has a recursive *revises* relationship, which links a new result obtained through reassessment to the previous damage assessment result; this relationship enables the revision history to be maintained and traced. The *revises* relationship is created when the application-level workflow explicitly designates a new assessment result for the same component as a reassessment or update of an existing *DamageResult*.

As an exception, *DamageResult* directly includes *assessment identifier* and *component identifier* in the relational-spatial data model. This is because the assessment result is the final damage assessment output at the building-component level, and therefore needs to support the direct retrieval of damage grades and spatial visualization at the corresponding object level.

3. Implementation

The proposed data model was implemented using a post-hurricane roof-damage assessment case. The implementation demonstrates how roof-damage assessment information can be organized and queried across spatial and graph databases. Multi-source aerial observations and field-based damage information were used to construct the database objects and relationships required for explainable assessment.

3.1. Study Area

The study area is a coastal residential area in Southern Texas, USA, affected by Hurricane Harvey, which made landfall along the Texas coast as a Category 4 hurricane on 25 August 2017. The middle Texas coastal region was strongly affected during 25–26 August 2017, exposing residential buildings to severe hurricane-related impacts. In such post-hurricane environments, roof damage may appear in various forms, including surface disturbance, debris accumulation, partial structural loss, and roof collapse. These damage patterns are not always represented by a single categorical label, because optical appearance and structural surface conditions may change in different ways depending on the type and severity of damage.

The selected area mainly consists of low-rise residential buildings, making it suitable for object-level and component-level roof-damage assessment. From a data-modeling perspective, the site provides a useful implementation case because roof damage can be described through multiple types of evidence. Optical imagery captures visual changes such as altered color, texture, edge patterns, exposed materials, and debris, whereas post-disaster surface-height data captures structural cues such as roof-surface loss, low-height flattening, height depression, surface roughness, and fragmentation. Therefore, the study area supports the construction of an evidence-aware data model in which building objects,

roof components, evidence records, damage indices, assessment results, and inferred change events can be explicitly linked.

3.2. Data

The implementation used multi-source aerial observations acquired before and after Hurricane Harvey. The pre-disaster optical observation was a National Agriculture Imagery Program (NAIP) aerial orthophoto acquired on 30 September 2016. The post-disaster observations consisted of a UAV orthomosaic and a UAV-derived DSM acquired on 9 September 2017. These datasets provide complementary information for roof-damage evidence generation. The pre- and post-disaster optical images support object-level optical evidence, while the post-disaster DSM provides the basis for surface-height evidence.

Figure 2 shows the main data layers used in this study. NAIP orthophoto was used as the reference image for building-footprint delineation and pre-event optical observation. The post-disaster UAV orthomosaic provided high-resolution optical information after the hurricane, while the UAV-derived DSM was converted into nDSM-like relative surface-height data by removing local ground elevation. These optical and surface-height datasets provided the observational basis for generating roof-damage evidence at the object level.




Figure 2. Multi-source data used for implementing the proposed data model. (a) Pre-disaster NAIP orthophoto, (b) post-disaster UAV orthomosaic, (c) DSM-derived relative surface-height data.

Building footprints were manually delineated from the pre-disaster NAIP orthophoto and used as object-level spatial units. Each building object was linked to a roof component so that the assessment could be represented at the building-component level. For each component, the inner evidence region defined in Section 2.1 was derived from the corresponding building footprint and used as the common spatial support for evidence extraction.

Field-based damage information was obtained from the Hurricane Harvey reconnaissance dataset [26], with matched records available for 80 of the 225 assessment units. The dataset provides separate component-level damage observations for roof cover, roof sheathing, and roof framing, in addition to an overall building damage rating. Because the overall building damage rating may also reflect damage to non-roof components such

as walls, windows, and doors, it was retained only as supplementary information and was not used as the direct reference for the present roof-focused assessment. Instead, the roof-cover, roof-sheathing, and roof-framing damage percentages were normalized to [0,1] and averaged to construct a roof-focused field damage indicator. The resulting indicator was then grouped into five ordinal roof-reference grades representing increasing levels of roof damage. These field reference grades were compared on an object-by-object basis with the independently derived C0–C4 assessment grades for the 80 matched records, and the proportion of matching grades was used as the assessment-level overall accuracy A_{OA} . The field reference was not used to determine or calibrate the C0–C4 decision boundaries, which were derived independently from the empirical distribution of I_i using the Jenks natural breaks procedure described in Section 3.3.1. Accordingly, the availability of a matched field reference does not imply independent validation of the corresponding C0–C4 grade.

3.3. Database Construction

3.3.1. Assessment Result and Rule Instantiation

Before constructing the hybrid database, each building footprint was represented as a Building object and linked to a roof BuildingComponent. The six evidence quantities defined in Section 2.1 were generated from the corresponding surface-height and optical observations and stored as DamageEvidence records. These evidence values were integrated into the continuous roof-damage evidence index I_i , which was stored as DamageIndex.

The empirical distribution of I_i across the 225 assessment units was partitioned into five groups using Jenks natural breaks [27]. The resulting boundaries were represented at two-decimal precision as 0.18, 0.32, 0.55, and 0.78 (Table 4), and the baseline damage grades were defined as

$$D_i = \begin{cases} C0, & 0.00 \leq I_i < 0.18 \\ C1, & 0.18 \leq I_i < 0.32 \\ C2, & 0.32 \leq I_i < 0.55 \\ C3, & 0.55 \leq I_i < 0.78 \\ C4, & 0.78 \leq I_i \leq 1.00 \end{cases} \quad (12)$$

C0–C4 represent increasing levels of roof-damage evidence, from little or no evidence to evidence consistent with severe roof damage. These boundaries are case-specific decision parameters rather than universally calibrated structural-damage thresholds.

Table 4. Jenks-based decision boundaries used for the case-study rule instantiation.

Boundary	Class Separation	Operational Threshold	Grade Boundary
T_1	0.1802	0.18	C0–C1
T_2	0.3159	0.32	C1–C2
T_3	0.5505	0.55	C2–C3
T_4	0.7789	0.78	C3–C4

In addition to the baseline index-based grading, two surface-evidence override rules were defined to represent extreme geometric evidence:

$$0.90 \leq S_i \leq 1.00, \quad (13)$$

and

$$0.00 \leq S_i \leq 0.10 \quad (14)$$

respectively. The high condition overrides the baseline index grade and assigns C4 with the component-specific change hypothesis RoofCollapseCandidate and needs_manual_review.

The low condition assigns C0 with NoChange and no_update_required. The latter is an operational rule and does not imply the absence of all possible non-geometric roof damage.

Assessment-level reliability was represented separately from the C0–C4 grade boundaries. Here, A_{OA} denotes the overall accuracy between the five-level roof-focused field reference and the C0–C4 assessment grades for the 80 matched field records. In the present case-study implementation, the reliability thresholds of 0.50 and 0.80 were introduced as operational criteria rather than empirically optimized or universally applicable thresholds. The 0.50 threshold was used to identify assessments with insufficient reliability, whereas the 0.80 threshold was used as a stricter criterion for downstream update eligibility. These thresholds are configurable rule parameters and may be adjusted for other assessment workflows. Results at or below the predefined reliability criterion were assigned low_confidence. For eligible index-based C4 results, specifically those produced by the baseline rule RUL_CV_C4, a stricter reliability condition was used as an operational gate for ready_to_update:

$$\text{RUL_CV_C4 satisfied} \wedge 0.80 \leq A_{OA} \leq 1.00 \quad (15)$$

The resulting assessment-level reliability value of 0.813 was retained in DamageAssessment as reliability metadata and was not used to determine or calibrate the Jenks-based C0–C4 boundaries.

The complete decision-rule configuration is summarized in Table 5. Each AssessmentRule includes a priority value that determines the evaluation order, with lower values applied first. Reliability control is evaluated first, followed by the surface-evidence override rules and the mutually exclusive index-based grading rules; the update readiness gate is subsequently applied to eligible index-based C4 results.

This explicit separation of evidence, index values, rule conditions, actions, and rule priority preserves the decision logic independently of the continuous evidence index and enables the applied rule to be traced for each damage result. The local stability of the case-study rule configuration is examined in Section 3.3.2.

3.3.2. Threshold and Index Sensitivity Analysis

Sensitivity analyses were conducted to assess the local robustness of the case-study rule configuration to variations in the index-based decision boundaries and continuous evidence index, rather than to optimize or recalibrate the decision parameters. Stability was evaluated using exact grade agreement, quadratic weighted kappa (QWK), the proportion of assessment units shifting by two or more grades, and the agreement for the combined C3–C4 group.

First, the sensitivity of the four index-based decision boundaries defined in Section 3.3.1 was evaluated. Let τ_j denote the j -th decision boundary, where $j \in \{1, 2, 3, 4\}$. The baseline threshold vector was defined as

$$\tau = (0.18, 0.32, 0.55, 0.78). \quad (16)$$

In a one-at-a-time analysis, each boundary was independently perturbed by $\pm 5\%$ and $\pm 10\%$, while the remaining boundaries were held fixed:

$$\tau'_j = \tau_j(1 + \delta), \quad (17)$$

where δ represents the specified perturbation rate. The complete decision procedure in Section 3.3.1 was reapplied for each perturbed configuration. Across all scenarios, fewer than 5% of assessment units changed grade, with a maximum change rate of 4.4% under the $\pm 10\%$ perturbation. All changes occurred between adjacent grades and were concentrated near the corresponding decision boundaries.

Table 5. Decision rules and precedence used in the case-study implementation.

Priority	Rule ID	Rule	Condition	Grade Action	Inferred Change	Update Readiness
1	RUL_ERR_LOW	Assessment Reliability Control	$0.00 \leq A_{OA} \leq 0.50$	Applies to all provisional C0–C4 grades	Unknown	low_confidence
2	RUL_EXC_ABO	High Surface-Evidence Override	$0.90 \leq S_i \leq 1.00$	C4	RoofCollapseCandidate	needs_manual_review
2	RUL_EXC_CON	Low Surface-Evidence Override	$0.00 \leq S_i \leq 0.10$	C0	NoChange	no_update_required
3	RUL_CV_C0	Grade C0 Baseline	$0.00 \leq I_i < 0.18$	C0	NoChange	no_update_required
3	RUL_CV_C1	Grade C1 Low	$0.18 \leq I_i < 0.32$	C1	MinorDamage	needs_additional_evidence
3	RUL_CV_C2	Grade C2 Moderate	$0.32 \leq I_i < 0.55$	C2	MinorDamage	needs_additional_evidence
3	RUL_CV_C3	Grade C3 High	$0.55 \leq I_i < 0.78$	C3	MajorDamage	needs_manual_review
3	RUL_CV_C4	Grade C4 Severe	$0.78 \leq I_i \leq 1.00$	C4	SevereDamage	needs_manual_review
4	RUL_UPDATE_CONF	Assessment Reliability Update Gate	RUL_CV_C4 satisfied $\wedge 0.80 \leq A_{OA} \leq 1.00$	C4	SevereDamage	ready_to_update

Simultaneous perturbation of all four boundaries was evaluated through 5000 Monte Carlo realizations:

$$\tau_j^{(r)} = \tau_j(1 + \delta_j^{(r)}), \quad (18)$$

where $\delta_j^{(r)}$ was independently sampled within either $\pm 5\%$ or $\pm 10\%$ of the corresponding baseline threshold. Mean exact grade agreement was 96.7% and 93.0%, with QWK values of 0.991 and 0.980, respectively. Agreement for the combined C3–C4 group remained 98.7% and 97.6%. Sensitivity to small variations in the continuous evidence index was evaluated independently as

$$I_i^i = \text{clip}(I_i + \varepsilon_i, 0, 1), \quad (19)$$

with the perturbation term defined as

$$\varepsilon_i \sim \mathcal{U}(-d, d), d \in \{0.01, 0.02\} \quad (20)$$

For perturbation ranges of ± 0.01 and ± 0.02 , 5000 realizations were generated. Exact grade agreement was 98.2% and 95.5%, with QWK values of 0.995 and 0.987, respectively. C3–C4 group agreement remained 99.3% and 98.9%, respectively.

The two surface-evidence override thresholds $S_i = 0.10$ and $S_i = 0.90$ were also independently perturbed by $\pm 5\%$ and $\pm 10\%$. Exact agreement ranged from 98.67% to 100.00%, with QWK values from 0.9963 to 1.0000, and no assessment unit shifted by two or more grades.

Stability was evaluated using exact grade agreement, QWK, the proportion of assessment units shifting by two or more grades, and the agreement for the combined C3–C4 group. The principal results are summarized in Table 6.

Table 6. Sensitivity of categorical damage grade assignments to perturbations in the decision boundaries and continuous evidence index.

Sensitivity Test	Perturbation	Exact Grade Agreement	QWK	≥ 2 -Grade Shift	C3–C4 Group Agreement
Decision boundaries	$\pm 5\%$	96.7%	0.991	0.0%	98.7%
Decision boundaries	$\pm 10\%$	93.0%	0.980	0.0%	97.6%
Continuous index I_i	± 0.01	98.2%	0.995	0.0%	99.3%
Continuous index I_i	± 0.02	95.5%	0.987	0.0%	98.9%

Overall, the categorical assignments were relatively stable under the tested perturbations, with changes largely limited to assessment units located near the corresponding decision boundaries. These results indicate local robustness of the case-study-specific rule configuration and should not be interpreted as evidence that the selected thresholds are universally optimal or transferable to other disasters, sensors, or assessment workflows. The assessment-level reliability thresholds of 0.50 and 0.80 were not included because they control reliability and update readiness states rather than the C0–C4 grade boundaries.

3.3.3. Hybrid Database Construction

To validate the effectiveness of the proposed model, a hybrid database was constructed using the damage assessment results described in Section 3.3.1. The implementation was carried out using two database systems. The model entities listed in Table 2 were implemented as a relational-spatial database in PostGIS, whereas the relationships defined in Table 3 were implemented as a graph database in Neo4j. This structure was adopted to support both the direct retrieval of damage assessment results and the semantic tracing of the assessment process through graph relationships.

In the PostGIS database, disaster information, damage assessment records, including damage evidence, damage index, damage result, and assessment rules, and inferred change events were stored as relational tables according to the entity definitions presented in Table 2. Buildings and building components were stored with their geometries and attribute information, as they represent the spatial objects which are assessment targets. Each table was designed with a unique identifier. Through these identifiers, damage assessment information, including building- and component-level damage results, can be directly retrieved from the spatial database and visualized with the corresponding spatial objects.

In the Neo4j database, 12 of 13 labeled relationships defined in Table 3 were instantiated. The *revises* relationship had no instance in the constructed database because no reassessment was included. The graph was designed to represent the semantic paths required to explain how each damage result was derived. Specifically, the graph structure enables the tracing of the damage index and assessment rule used to determine a given damage result, as well as the damage evidence from which the corresponding index was computed. In addition, the explicit relationships among buildings, building components, damage results, and inferred change events were constructed to ensure the explainability of assessment results.

Consequently, 12 relational tables were generated in the PostGIS database with spatial indices. In the Neo4j database, the instances of 12 entity types were constructed as 2510 nodes, and 6999 relationships were generated according to the predefined relationship types. Data ingestion into PostGIS required 0.058 s, while the construction of 2510 nodes and 6999 relationships in Neo4j required 1.416 s. The resulting database sizes were approximately 1.14 MB for the implemented PostGIS model, including the supporting reference tables, and 2.24 MB for the Neo4j database store. The hybrid database based on the proposed data model was implemented to support both spatial retrieval of damage assessment results and explanation queries based on graph traversal.

4. Case Study

4.1. Competency Questions for Explainable Assessment

The competency questions (CQs) were used to validate whether the proposed model can retrieve and return information required for explainable disaster damage assessment. Specifically, the CQ-based validation evaluates the retrievability and completeness of the traceable rule-based explanations supported by the proposed data model. Defining a set of CQs is an important step in data modeling, and the resulting questions can also be used to evaluate the schema derived from the modeling process [28]. In this study, we designed six different CQs from the intended information requirements of the proposed assessment workflow and hybrid data model to validate three aspects of the data model, as shown in Table 7. Table 8 presents success criteria and validation scope for each CQ. First, we examined whether the hybrid database constructed based on the proposed data model appropriately reflects the workflow of building damage assessment in the disaster domain. In disaster damage assessment, once data for the target area are provided as input, damage evidence observed from the data is generated, and assessment indices are calculated through combinations of such evidence. By applying appropriate rules to the damage assessment evidence and indices, assessment results are derived, and building change events are inferred from these results. CQ1 and CQ2 examine whether damage assessment information and logic can be properly retrieved along this flow.

Next, CQ3 and CQ4 were designed to validate the explainability requirement, which constitutes the core contribution of the proposed data model. This is directly related to the necessity of the proposed data model. Rather than simply presenting assessment results as numerical index values, these CQs examine whether it is possible to generate human-interpretable explanations of which index and which rule were used to derive a specific

damage assessment result, how that result led to the inference of a change event, and from which evidence a specific index was derived.

Table 7. Design of competency questions.

No	Validation Dimension	Representative CQ	Query Type
1	Reference assessment flow	What is the current damage result of a specific building component?	Table query
2	Reference assessment flow	What change event and update readiness status were inferred from a specific damage result?	Multi-hop graph traversal
3	Explainability, Provenance	Which damage index and evidence support a specific damage result, and from which assessment inputs were the evidence derived?	Multi-hop graph traversal
4	Explainability, Rule traceability	Which assessment rule was applied to derive a specific damage result?	Multi-hop graph traversal
5	Spatial retrieval	Where are building components with current severe-damage results located, and which damage results are associated with them?	Spatial query and map visualization
6	spatial-graph integration	For a selected damaged component on the map, what current result, damage index, evidence and assessment inputs, applied rule, and change event explain its assessment?	Hybrid cross-database query

Table 8. Success criteria and validation scope for the competency questions.

CQ	Success Criteria	Validation Scope
CQ1	Exactly one current result per component.	All 225 building components
CQ2	Exactly one inferred change event and update readiness status per result.	All 225 damage results
CQ3	Complete result–index–evidence–input provenance.	All 225 damage results
CQ4	Complete result–rule–condition–action traceability.	All 225 damage results
CQ5	Complete spatial retrieval of current C4 results.	All 16 current C4 damage results
CQ6	Complete and consistent cross-database explanation linking a selected component to all associated assessment information.	All 225 component–result pairs

Finally, we aimed to verify the suitability of the hybrid database implementation and the intended use of the proposed data model, thereby confirming the validity of the integrated use of a relational-spatial database and a graph database. CQ5 and CQ6 were designed to validate the integrated use of PostGIS and Neo4j databases for spatially retrieving and semantically interpreting assessment results. Specifically, they examine whether geometry-based spatial queries can be performed at the building-component level in PostGIS and whether assessment results can be retrieved at the object level. In addition, they examine whether Neo4j can support multi-hop explainability across the evidence–index–rule–result–change path.

4.2. Query Results

For each CQ, the natural language query was first automatically converted into SQL and Cypher queries using an LLM (GPT-5.5). The LLM was used only as an auxiliary query-drafting tool and was not a component of the proposed framework. The generated queries were then manually reviewed and executed to return the corresponding results. The

final SQL and Cypher queries used in the CQ validation are provided in the Supplementary Materials. The query execution results are summarized in Table 9. Execution times are reported as the median of ten measured runs after three warm-up runs, with database connection time excluded. For CQ6, the execution time represents the end-to-end retrieval from the PostGIS lookup to the corresponding Neo4j explanation.

For CQ1 and CQ2, it was confirmed that the required information could be appropriately retrieved according to the damage assessment flow. CQ1, which corresponds to a simple table retrieval query, properly returned the damage assessment result of Roof-Damage with damage grade C2 for the input component COMP_BLD_000000_roof. In CQ2, the inferred change event derived from a specific assessment result was successfully returned by querying the [InferredChangeEvent–inferred_by -> DamageResult] path in the Neo4j database. From the damage assessment result RES_BLD_000001, a building change event corresponding to RoofCollapseCandidate was inferred, and the result further indicated that manual review is required before any downstream database-update decision. CQ1 confirmed that the proposed data model can directly retrieve damage assessment information at the spatial object level, while CQ2 showed that the proposed model can appropriately link the final damage assessment result to a potential post-disaster change event and additionally provide information related to database updating.

In CQ3 and CQ4, the supporting evidence and decision rules associated with a damage assessment result were traced and returned according to the queries as shown in Figure 3. For the RES_BLD_000001, graph traversal returned the related damage index and the six evidence elements used to compute that index (Figure 3a). In addition, it was possible to retrieve that the damage grade C4 was derived from the *High Surface-Evidence Override* rule (Figure 3b).




Figure 3. Graph traversal results for (a) CQ3 and (b) CQ4.

As shown in Table 10, by defining an appropriate natural language pattern and using the returned values, an explanation of the assessment result could be generated; the generated explanation also included an IF–THEN expression for the rule. For example, the explanation pattern for CQ4 was formulated as follows: The damage assessment result {result_id} was derived from the assessment rule {rule_name}. The rule was reconstructed as an IF–THEN expression: IF {rule_condition} THEN {rule_action}. In other words, CQ3 and CQ4 confirmed that the proposed model can explicitly represent the evidence–index–result path through the separately defined semantic relationship graph. They also showed that the model can support intuitively understandable explanations by retrieving and reconstructing the rule-based logic used to derive the assessment result.

Table 9. Summary of competency question results.

CQ	Query Environment	Result and Execution Time (ms)	Representative Parameter	Representative Result Summary
CQ1	PostGIS	225/225 passed; 0.096	component_id = COMP_BLD_000000_roof	damage type = RoofDamage; damage grade = C2
CQ2	Neo4j	225/225 passed; 2.573	result_id = RES_BLD_000001	change type = RoofCollapseCandidate; update readiness = needs_manual_review
CQ3	Neo4j	225/225 passed; 2.985	result_id = RES_BLD_000001	Result RES_BLD_000001 was traced to its damage index, six supporting evidence records, and their source assessment inputs.
CQ4	Neo4j	225/225 passed; 3.688	result_id = RES_BLD_000001	Result RES_BLD_000001 was traced to the applied rule RUL_EXC_ABO, including its condition and action.
CQ5	PostGIS	16/16 passed; 0.179	Damage-grade = C4	16 severely damaged building component(s) were retrieved and joined with geometries for map visualization.
CQ6	PostGIS + Neo4j	225/225 passed; 3.806	component_id = COMP_BLD_000020_roof	Result RES_BLD_000020 was retrieved with its building geometry in PostGIS and linked graph explanation in Neo4j, including I_i , supporting evidence, applied rule, and inferred change event.

Table 10. Generated explanations for CQ3 and CQ4.

CQ	Generated Explanation
CQ3	The damage assessment result RES_BLD_000001 for building BLD_000001 was classified as RoofDamage with damage grade C4. The damage index $I_i = 1.0$. was computed from the following evidence: Surface evidences $H_i = 1$, $G_i = 0$, $S_i = 1$; Optical evidences $O_i^{post} = 0.516$; $\Delta O_i = 0.513$; $O_i = 0.265$. The surface evidence was derived from the post-event surface input (UAV-derived DSM), O_i^{post} from the post-event optical input (UAV orthomosaic), and ΔO_i and O_i from both the pre-event optical input (NAIP aerial orthophoto) and the post-event optical input (UAV orthomosaic).
CQ4	The damage assessment result RES_BLD_000001, classified as RoofDamage with damage grade C4, was derived from the assessment rule <i>High Surface-Evidence Override (RUL_EXC_ABO)</i> . The rule was reconstructed as an IF-THEN expression: IF $0.9 \leq S_i \leq 1.0$ THEN damage_grade = C4 AND change_type = RoofCollapseCandidate AND update_readiness = needs_manual_review. For this result, $S_i = 1.0$, satisfying the rule condition.

Figure 4 shows the visual explanation for RES_BLD_000001 by linking the input observations and the evidence records used in the generated graph-based explanation. Figure 4a,b shows the pre-event NAIP orthophoto and post-disaster UAV orthomosaic used for optical evidence extraction, while Figure 4c shows the post-disaster relative surface-height data used for surface-evidence extraction. The building footprint represents the pre-event object boundary, and slight apparent boundary offsets may remain because the pre- and post-disaster datasets differ in spatial resolution and imaging conditions, especially for small buildings. Therefore, the dashed inner evidence region was used as the spatial support for evidence extraction to reduce boundary effects and residual misregistration. The surface and optical evidence values visualized in Figure 4d are identical to those returned by the graph query and used in the generated explanation. This consistency indicates that the query-derived explanation properly reflects the evidence extracted from the input data and that the proposed model can support traceable interpretation from input observations to damage assessment results.

In CQ5, the spatial locations of buildings classified into the most severe-damage grade were retrieved, and the query results were visualized on a map. As shown in Figure 5a, 16 records classified as C4 were successfully returned. Using the geometries returned together with the query results, the building shapes were visualized at their correct locations on the map while preserving their geometries. This result demonstrates that the proposed model can directly retrieve spatial objects associated with assessment results and represent their spatial distribution using the stored geometries and coordinates. Figure 5b shows that the retrieved buildings exhibit clear visual evidence of severe roof damage, such as roof-surface loss, exposed structural elements, and debris within the building footprints. This result demonstrates that the hybrid query can retrieve severe-damage cases at the study-area scale while maintaining a direct link to object-level visual evidence.

In CQ6, for the specific damage assessment result RES_BLD_000020, the related building geometry and centroid coordinates were returned through a PostGIS spatial query, while an explanation graph composed of six semantic relationships was returned through Neo4j graph traversal. RES_BLD_000020 was returned as the assessment result for building BLD_000020 and component COMP_BLD_000020_roof. The spatial visualization generated using the returned centroid coordinates and geometry is shown in Figure 6.




Figure 4. Visual explanation for RES_BLD_000001: (a) pre-event NAIP orthophoto, (b) post-disaster UAV orthomosaic, (c) post-disaster relative surface-height data, and (d) evidence records supporting the roof-damage evidence index.




Figure 5. C4 damage results retrieved using CQ5: (a) spatial distribution map and (b) object-level image crops showing the corresponding building footprints and internal evidence regions.




Figure 6. Query result for CQ6 retrieved from the hybrid databases.

The Neo4j query result showed that the corresponding result was assigned damage grade C0. It was also confirmed that this result was derived by applying the Grade C0 Baseline rule to the damage index value computed from six evidence elements. In addition, the inferred change event derived from this result was NoChange, indicating that no update was required. The result of CQ6 confirmed that the hybrid database constructed based on the proposed model can support the spatial information of a specific building, its damage assessment information, and graph-based semantic explanation of the assessment result in an integrated manner.

Figure 7 shows the visual explanation for RES_BLD_000020 by linking the input observations and the evidence records returned by the hybrid query. As shown in Figure 7a–c, the target roof component within the building footprint shows no obvious roof-surface loss or collapse, although debris and damaged materials are visible in the surrounding area. This visual interpretation is consistent with the low surface and optical evidence values shown in Figure 7d. The evidence values presented in Figure 7d are identical to those included in the generated graph-based explanation. This consistency indicates that the query results reflect the evidence extracted from the input observations and that the proposed model can support traceable interpretation from object-level observations to damage assessment results.

Beyond the representative examples, each CQ was also evaluated over its full validation scope, including all 225 building components and all 16 current C4 results for CQ5. All cases satisfied the predefined success criteria, as shown in the *Result and Execution Time* column in Table 9. In addition to the CQ-based validation, four controlled edge-case tests were conducted to examine the proposed model under exceptional conditions. First, a missing-evidence case was simulated by temporarily removing one evidence record, and the provenance completeness check correctly detected the incomplete evidence path. Also, a multiple-rule-match case was tested using input values that simultaneously satisfied three rules; the *High Surface-Evidence Override* rule was selected according to the predefined rule priority. An out-of-range index value that satisfied no rule condition was introduced, and the model correctly identified the absence of a matching rule without assigning an arbitrary result. Finally, a repeated-assessment case was simulated by creating a new assessment and result for an existing component. The previous result was retained as historical information, only the new result remained current. During this test, the existing current DamageResult in PostGIS was closed by setting its *valid_to* value, and the new DamageResult was stored

as the current result within a single transaction. Correspondingly, the temporal state of the previous DamageResult was updated and a new DamageResult node was created in Neo4j, thereby maintaining synchronization between the two databases at the application level. Since the new result was explicitly treated as a reassessment, the revision history was additionally represented through the *revises* relationship.




Figure 7. Visual explanation for RES_BLD_000020. (a) Pre-event NAIP orthophoto, (b) post-disaster UAV orthomosaic, (c) post-disaster relative surface-height data, and (d) evidence records supporting the roof-damage evidence index.

5. Discussion

The proposed data model preserves the intermediate information used to derive a categorical damage result rather than storing only the final grade. In the present implementation, surface and optical evidence are represented separately, integrated into a continuous evidence index, and interpreted through explicit decision rules. This separation allows individual evidence variables to participate directly in the decision layer when required and enables the basis of an assigned grade to be reconstructed through the relationships among observations, evidence, index values, and applied rules.

The C0–C4 categories should be interpreted as evidence-based roof-damage grades for the present assessment workflow rather than as universal physical damage classes. The baseline boundaries were derived from the empirical distribution of the continuous evidence index using Jenks natural breaks, while extreme surface-evidence conditions were represented separately as explicit rule conditions. Sensitivity analysis showed that the categorical assignments remained relatively stable under the tested variations in the index-based boundaries and continuous evidence values. However, this local stability does not imply transferability of the numerical thresholds to other disasters, sensors, or assessment

workflows; these thresholds should be treated as case-specific decision parameters that may require re-estimation or adjustment.

Building footprints were manually delineated to provide a consistent spatial reference and to avoid introducing additional segmentation uncertainty into the present evaluation. In operational applications, existing building-footprint databases or automated extraction methods may instead be used, although temporal mismatch, positional errors, and boundary inaccuracies can propagate into object-level evidence extraction. Because the pre-disaster footprint is retained as the spatial reference, evidence can still be extracted when the post-disaster roof boundary becomes visually indistinct. The inner evidence region reduces boundary contamination and residual misregistration, but attribution uncertainty may remain when debris from adjacent structures enters the target region. The quantitative effects of automated segmentation errors were not evaluated and remain to be examined in future operational implementations.

The pre- and post-disaster optical observations also differ in acquisition date, platform, spatial resolution, and imaging conditions. The present workflow therefore uses object-level optical comparison rather than direct pixel-to-pixel differencing and combines optical information with surface evidence. Nevertheless, non-disaster changes occurring between acquisition dates, such as roof renovation or recoating, cannot be completely excluded. The data-model structure itself is not tied to a specific sensor or UAV platform, but the evidence, index, and rule configuration must correspond to the observations actually available. The present case-study configuration requires both surface and optical evidence and should not be applied unchanged when a required evidence source is unavailable; as defined in Section 2.1, missing evidence is not represented by assigning a zero value. Alternative observation configurations remain to be evaluated.

Field reconnaissance records were matched to 80 of the 225 assessment units and included both overall building damage and roof-component observations. The roof-related fields were used to construct the roof-focused reference for the assessment-level evaluation reported in Section 3.3.1, whereas the overall building rating may also reflect damage to components not observable from the roof-oriented aerial assessment. Importantly, the field observations were not used to determine or calibrate the Jenks-based C0–C4 boundaries. The resulting assessment-level reliability information should therefore be distinguished from object-level validation: the existence of a matched field record does not by itself establish the correctness of the corresponding C0–C4 grade or inferred change.

Finally, the present case study empirically demonstrates the proposed structure only for hurricane-induced roof damage, and the implementation does not explicitly define vertical boundaries among roof, facade, and lower-floor components. The current evidence variables, index formulation, and numerical thresholds should therefore not be assumed to generalize directly to other components or disaster types. For a different assessment workflow, the core entity and relationship structure can be retained and instantiated according to the target component, available observations, and assessment procedure.

6. Conclusions

In this study, an evidence-aware data model was proposed to make the post-disaster building damage assessment flow reusable and to manage assessment results as traceable information. For damage assessment results to be practically used for post-disaster change inference, spatial database updating, and other applications, it is necessary to go beyond storing assessment results simply as numerical values or final grades. Instead, the evidence, indicators, rules, and results generated during the assessment process should be connected within a single structured workflow. In this respect, this study formalized the process by which damage assessment results are derived and the logic by which they are interpreted,

and implemented the model as a hybrid-structured database combining a spatial database and a graph database.

The proposed model was designed based on a component-level roof-damage assessment workflow using optical imagery and post-disaster surface-height data. Not only the surface- and optical-derived evidence, but also the damage indices calculated from combinations of evidence were defined as explicit entities. Furthermore, the model encompasses the final assessment results as well as the decision rules used to derive those results. In addition, the proposed data model was implemented as a hybrid database consisting of a PostGIS database and a Neo4j graph database, so that spatial object management and multi-hop relationship tracing could be supported in a complementary manner.

The validity of the proposed model was evaluated through six CQs derived from various requirements. The results confirmed that the proposed model can support the object-level retrieval of damage results, inference of change events caused by those results, and tracing of the explanation path of assessment results that connects damage evidence, damage indices, and assessment rules. In addition, through spatial queries, the geometries and attributes of damaged buildings could be retrieved and visualized on a map. It was also demonstrated that the spatial information of a specific building, its damage assessment result and evidence, and the inferred change event could be connected within a single traceable structure.

In addition, the sensitivity analysis showed that the case-study categorical assignments remained largely stable under moderate perturbations of the principal decision boundaries and continuous evidence values.

This study does not aim to develop a new damage assessment framework or to improve the performance of automated damage-detection algorithms. Instead, it focuses on structuring relevant information with reference to a predefined workflow and managing it in a traceable manner. In the present implementation, the damage evidence and indices were defined based on optical imagery and post-disaster surface-height data. When applying the model to other disaster types, sensor configurations, or damage patterns, the definitions of evidence indicators, damage indices, and rule thresholds should therefore be re-examined and adjusted. Future research needs to extend the model so that it can accommodate various observation data, damage indices, and assessment rules, and to verify its operational applicability by linking it with actual damage assessment systems or spatial database updating workflows. Although the proposed data model was designed for extensibility, the present implementation and case-study evaluation are limited to hurricane-induced roof damage. Therefore, its applicability to other disaster types, building components, and damage patterns should be evaluated in future studies.

The proposed hybrid database approach may cause additional operational complexity compared to a single-store architecture. Although we evaluated case-study-scale data ingestion time, graph construction time, database size, and end-to-end query latency, these measurements should not be interpreted as a production-scale operational evaluation. Key operational issues, including consistent updates across the two databases, synchronization latency, and separate backup strategies, remain to be addressed in future work.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/ijgi15090414/s1>, S1: Schema Definition; S2: SQL and Cypher Queries for CQ 1-6.

Author Contributions: Conceptualization, Ahram Song and Seula Park; methodology, Ahram Song and Seula Park; validation, Ahram Song and Seula Park; formal analysis, Ahram Song and Seula Park; resources, Jinha Jung and Anjin Chang; data curation, Anjin Chang, Ahram Song and Seula

Park; writing—original draft preparation, Ahram Song and Seula Park; writing—review and editing, all authors. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2025-24534253).

Data Availability Statement: Restrictions apply to the datasets used in this study. The post-event UAV orthomosaic, UAV-derived DSM, manually delineated building footprints, and derived geospatial/database products are not publicly available due to data-use and redistribution restrictions.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (GPT-5.6, OpenAI) to assist with translation for improving readability and with the refinement of Python scripts used for data processing. The authors reviewed, edited, and validated the outputs and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Shin, D.Y.; Kim, D.J.S.; Kim, S.S.; Han, Y.K.; Nho, H.J. A Study on the Use of Drones for Disaster Damage Investigation in Mountainous Terrain. *Korean J. Remote Sens.* **2020**, *36*, 1209–1220. [\[CrossRef\]](#)
2. Ser, J.; Yang, B. Quantitative Evaluations of Deep Learning Models for Rapid Building Damage Detection in Disaster Areas. *J. Korean Soc. Surv. Geod. Photogramm. Cartogr.* **2022**, *40*, 381–391. [\[CrossRef\]](#)
3. Gupta, R.; Goodman, B.; Patel, N.; Hosfelt, R.; Sajeev, S.; Heim, E.; Doshi, J.; Lucas, K.; Choset, H.; Gaston, M. Creating xBD: A Dataset for Assessing Building Damage from Satellite Imagery. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, Long Beach, CA, USA, 16–21 June 2019; pp. 10–17.
4. Hänsch, R.; Arndt, J.; Lunga, D.; Gibb, M.; Pedelose, T.; Boedihardjo, A.; Petrie, D.; Bacastow, T.M. Spacenet 8—The detection of flooded roads and buildings. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 19–20 June 2022; pp. 1472–1480. [\[CrossRef\]](#)
5. Chen, H.; Song, J.; Dietrich, O.; Broni-Bediako, C.; Xuan, W.; Wang, J.; Shao, X.; Wei, Y.; Xia, J.; Lan, C.; et al. BRIGHT: A Globally Distributed Multimodal Building Damage Assessment Dataset with Very-High-Resolution for All-Weather Disaster Response. *Earth Syst. Sci. Data* **2025**, *17*, 6217–6253. [\[CrossRef\]](#)
6. Cheng, C.S.; Behzadan, A.H.; Noshadravan, A. Deep learning for post-hurricane aerial damage assessment of buildings. *Comput.-Aided Civ. Infrastruct. Eng.* **2021**, *36*, 695–710. [\[CrossRef\]](#)
7. Khajwal, A.B.; Cheng, C.; Noshadravan, A. Post-disaster damage classification based on deep multi-view image fusion. *Comput.-Aided Civ. Infrastruct. Eng.* **2023**, *38*, 528–544. [\[CrossRef\]](#)
8. Fernandez Galarreta, J.; Kerle, N.; Gerke, M. UAV-based urban structural damage assessment using object-based image analysis and semantic reasoning. *Nat. Hazards Earth Syst. Sci.* **2015**, *15*, 1087–1101. [\[CrossRef\]](#)
9. Kerle, N.; Nex, F.; Gerke, M.; Duarte, D.; Vetrivel, A. UAV-Based Structural Damage Mapping: A Review. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 14. [\[CrossRef\]](#)
10. Zhou, Z.; Gong, J.; Hu, X. Community-Scale Multi-Level Post-Hurricane Damage Assessment of Residential Buildings Using Multi-Temporal Airborne LiDAR Data. *Autom. Constr.* **2019**, *98*, 30–45. [\[CrossRef\]](#)
11. Bhanumurthy, V.; Ram Mohan Rao, K.; Jai Sankar, G.; Nagamani, P.V. Spatial Data Integration for Disaster/Emergency Management: An Indian Experience. *Spat. Inf. Res.* **2017**, *25*, 303–314. [\[CrossRef\]](#)
12. Ghaffarian, S.; Kerle, N.; Pasolli, E.; Arsanjani, J.J. Post-Disaster Building Database Updating Using Automated Deep Learning: An Integration of Pre-Disaster OpenStreetMap and Multi-Temporal Satellite Data. *Remote Sens.* **2019**, *11*, 2427. [\[CrossRef\]](#)
13. Amirebrahimi, S.; Rajabifard, A.; Mendis, P.; Ngo, T. A data model for integrating GIS and BIM for assessment and 3D visualisation of flood damage to building. *Locate* **2015**, *1323*, 78–89.
14. Amirebrahimi, S.; Rajabifard, A.; Mendis, P.; Ngo, T. A BIM-GIS integration method in support of the assessment and 3D visualisation of flood damage to a building. *J. Spat. Sci.* **2016**, *61*, 317–350. [\[CrossRef\]](#)
15. Biljecki, F.; Stoter, J.; Ledoux, H.; Zlatanova, S.; Çöltekin, A. Applications of 3D city models: State of the art review. *ISPRS Int. J. Geo-Inf.* **2015**, *4*, 2842–2889. [\[CrossRef\]](#)
16. Biljecki, F.; Kumar, K.; Nagel, C. CityGML Application Domain Extension (ADE): Overview of developments. *Open Geospat. Data Softw. Stand.* **2018**, *3*, 13. [\[CrossRef\]](#)
17. Shen, J.; Zhou, J.; Zhou, J.; Herman, L.; Reznik, T. Constructing the CityGML ADE for the Multi-Source Data Integration of Urban Flooding. *ISPRS Int. J. Geo-Inf.* **2020**, *9*, 359. [\[CrossRef\]](#)
18. Luo, S.; Ding, L.; Fan, H. A CityGML ADE for flood simulations. *Geo-Spat. Inf. Sci.* **2026**, *29*, 2438–2452. [\[CrossRef\]](#)

19. Chaidas, K.; Tataris, G.; Soulakellis, N. Seismic Damage Semantics on Post-Earthquake LOD3 Building Models Generated by UAS. *ISPRS Int. J. Geo-Inf.* **2021**, *10*, 345. [[CrossRef](#)]
20. Open Geospatial Consortium. *GeoSPARQL: A Geographic Query Language for RDF Data*; OGC: Arlington, VA, USA, 2023. Available online: <https://www.ogc.org/standards/geosparql> (accessed on 30 August 2026).
21. Lebo, T.; Sahoo, S.; McGuinness, D.; Belhajjame, K.; Cheney, J.; Corsar, D.; Garijo, D.; Soiland-Reyes, S.; Zednik, S.; Zhao, J. PROV-O: The PROV Ontology. 2013. Available online: <https://www.w3.org/TR/prov-o/> (accessed on 30 March 2026).
22. Haller, A.; Janowicz, K.; Cox, S.J.D.; Le Phuoc, D.; Taylor, K.; Lefrançois, M. *Semantic Sensor Network Ontology*; W3C Recommendation; World Wide Web Consortium: Wakefield, MA, USA, 2017. Available online: <https://www.w3.org/TR/vocab-ssn/> (accessed on 30 August 2026).
23. Arrieta, A.B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; Garcia, S.; Gil-Lopez, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. [[CrossRef](#)]
24. Hogan, A.; Blomqvist, E.; Cochez, M.; d’Amato, C.; Melo, G.D.; Gutierrez, C.; Kirrane, S.; Gayo, J.E.L.; Navigli, R.; Neumaier, S.; et al. Knowledge graphs. *ACM Comput. Surv.* **2021**, *54*, 71. [[CrossRef](#)]
25. Tiddi, I.; Schlobach, S. Knowledge Graphs as Tools for Explainable Machine Learning: A Survey. *Artif. Intell.* **2022**, *302*, 103627. [[CrossRef](#)]
26. Kijewski-Correa, T.; Gong, J.; Kennedy, A.; Womble, J.A.; Cai, S.; Cleary, J.; Dao, T.; Leite, F.; Liang, D.; Peterman, K.; et al. *RAPID: Structural Wind Engineering Reconnaissance of Hurricane Harvey: A Supplement to Collaborative Research: Geotechnical Extreme Events Reconnaissance (GEER) Association: Turning Disaster into Knowledge—Data Report*; NHERI DesignSafe: Austin, TX, USA, 2018.
27. Jenks, G.F. The data model concept in statistical mapping. *Int. Yearb. Cartogr.* **1967**, *7*, 186–190.
28. Rebboud, Y.; Tailhardat, L.; Lisena, P.; Troncy, R. Can LLMs Generate Competency Questions? In Proceedings of the European Semantic Web Conference, Hersonissos, Greece, 26–30 May 2025; pp. 71–80. [[CrossRef](#)]

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.