SICA: A Rubric-Based Vision-LLM Framework for Post-Earthquake Bridge Inspection

Issac K.T. Pang [(1)], Khalid M. Mosalam [(2)]

[(1)] *PhD Candidate, Dept. of Civil and Environmental Engineering, Univ. of California, Berkeley, issac.pang@berkeley.edu*

[(2)] *Taisei Professor of Civil Engineering and PEER Director, Dept. of Civil and Environmental Engineering, Univ. of California, Berkeley, mosalam@berkeley.edu*

Abstract

SICA (Structure Investigation and Condition Assessment) is a vision-enabled large language model (LLM) framework that supports rapid post-earthquake bridge inspection triage from photographs. The system evaluates inspection images by structural component and assigns a four-level, action-oriented triage rating to each photo with concise engineering reasoning. Photo-level outputs are aggregated to produce component- and event-level triage decisions for rapid operational prioritization.

SICA embeds explicit, rubric-based engineering decision criteria directly into the inference process, improving transparency and consistency relative to black-box image classification approaches and enabling engineers to review and customize assessment logic. The proposed framework is model-agnostic, enabling different vision-capable LLMs to be integrated without changes to the assessment pipeline, while producing consistent, structured outputs for downstream reporting and user interfaces. The framework is demonstrated through representative inspection scenarios, highlighting its ability to produce consistent, interpretable, and actionable triage decisions aligned with established engineering guidelines. To support practical deployment, SICA includes a mobile interface for in-field photo capture, structured data collection, and real-time assessment. This integration enables scalable, interpretable, and field-deployable post-earthquake bridge screening and decision support.

Keywords: *Decision support systems; Infrastructure resilience; Inspection triage; Post-earthquake bridge inspection; Vision-enabled large language models.*

1. Introduction

Rapid and reliable assessment of bridge conditions following extreme events, such as earthquakes, is critical for ensuring public safety and restoring transportation networks. The California Department of Transportation utilizes the First Responder Bridge Assessment Guide (FRBAG) [1] as a standardized protocol to support post-event inspection and decision-making. FRBAG provides engineers with visual criteria and rule-based guidance to determine whether a bridge should remain open or be closed, with or without further inspection. FRBAG functions as an action-oriented field resource that links observed damage conditions in key bridge components to immediate operational decisions and follow-up inspection needs. While effective, this process is inherently manual, relying heavily on human judgment, experience, and time-consuming field workflows.

Recent advances in computer vision [2-5] and large language models (LLMs) [6-8] present new opportunities to augment traditional inspection practices. In particular, vision-capable LLMs enable automated interpretation of visual data while retaining the ability to generate structured, human-readable reasoning. This creates the potential to translate domain-specific engineering guidelines, such as FRBAG, into scalable and consistent digital decision-support systems.

This paper presents SICA (Structure Investigation and Condition Assessment), a vision-based assessment framework designed to support post-earthquake bridge inspection using image data. SICA integrates inspection photos, structured engineering rubrics, and vision-capable LLMs to produce standardized condition classifications and accompanying explanations. In this framework, the action-oriented logic of FRBAG is formalized through an R-state classification (R1 to R4). These range from "open" to "close immediately," together with concise engineering reasoning. These results are further aggregated from the photo level to the component level, and ultimately to the event level, enabling rapid prioritization of inspections and response actions. A key feature of SICA is its transparent and modular design. Rather than embedding decision logic within model weights, SICA externalizes engineering knowledge through editable rubrics and prompt templates. This allows domain experts to review, modify, and extend assessment criteria without altering the core system architecture. In addition, the framework is model-agnostic, supporting integration with different vision models through a unified interface.

Moreover, this work introduces a mobile-centric inspection workflow, enabling field engineers to capture, organize, and analyze inspection data directly on-site. Compared to traditional desktop-based workflows, this approach reduces data handling overhead and supports near real-time decision-making in post-event scenarios.

The main contributions of this work are:

1) A rubric-based vision-LLM framework that explicitly embeds engineering decision logic into the assessment process.
2) A hierarchical triage methodology that aggregates photo-level assessments into component- and event-level decisions.
3) A mobile-centric workflow enabling real-time, in-field inspection and decision support.
4) A transparent and model-agnostic architecture that allows domain experts to modify assessment criteria without retraining models.

2. Overview of the SICA Framework

Fig. 1 illustrates the overall workflow of the SICA framework, which enables automated post-earthquake bridge condition assessment using inspection images.

The SICA framework takes as input a set of bridge inspection photographs, which are first organized according to key structural components, including approaches, columns, joints/hinges, and abutments. This component-based organization reflects standard engineering inspection practices and ensures that damage assessment is performed in a structured and interpretable manner. For each input image, SICA performs vision-based analysis using a vision-capable LLM. The system evaluates visible damage features and assigns a corresponding R-state classification (R1 to R4) [1], following the inspection logic of the FRBAG booklet. In general, R1 indicates no significant issues and supports normal operation, while R2 reflects minor concerns that warrant follow-up inspection while the bridge remains open. R3 corresponds to significant damage requiring closure or restricted access pending further inspection. R4, by contrast, denotes severe and unstable conditions in which the bridge is considered unsafe even for engineers to approach or inspect, due to the risk of imminent collapse. An R4 classification therefore requires immediate closure. In addition to the classification, SICA generates a *short engineering rationale* that explains the observed damage in terms consistent with domain-specific inspection criteria.

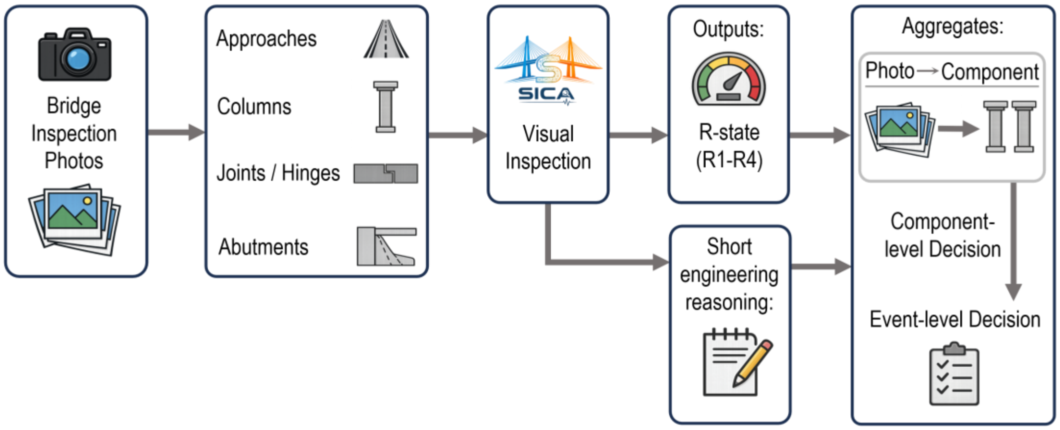

Fig. 1: Overview of the SICA framework for post-earthquake bridge condition assessment.

To support decision-making at different levels of granularity, SICA adopts a hierarchical aggregation strategy. First, image-level results are grouped by structural component, and a component-level decision is determined based on the most critical condition observed within that component. Subsequently, all component-level results are aggregated to produce an event-level decision, which reflects the overall safety condition of the bridge. This hierarchical design, from photo-level assessment to component-level evaluation and ultimately to event-level decision-making, ensures that critical damage is not overlooked while maintaining interpretability and traceability. By combining automated visual analysis with structured engineering reasoning, SICA provides a scalable and consistent framework for rapid post-earthquake bridge assessment. The mathematical foundation of this hierarchy starts by denoting $\mathcal{I}$ as the set of inspection images, grouped into a set of structural components $\mathcal{C} = \{1 \quad \cdots \quad C\}$. SICA defines a photo-level classifier:

$$f: \mathcal{I} \rightarrow Ri \in \{R1 \quad R2 \quad R3 \quad R4\} \tag{1}$$

For component-level decisions and for component $c \in \mathcal{C}$, define

$$Ri_c = \max_{i \in \mathcal{I}_c} f(i) \tag{2}$$

and the event-level decision is finally expressed as:

$$Ri_{\text{event}} = \max_{c \in \mathcal{C}} (Ri_c) \tag{3}$$

3. Rubric-Based Vision-LLM Assessment Method

The SICA framework adopts a rubric-based vision-LLM approach to enable structured and interpretable assessment of bridge inspection images. The key idea is to incorporate engineering knowledge into component-specific rubrics, rather than embedding decision logic implicitly within model parameters.

For each inspection image, the system constructs a prompt that integrates component context and rubric-defined damage criteria. These rubrics encode relevant damage indicators and their association with standardized condition states. A vision-capable LLM then processes the image together with this structured input to generate a consistent, schema-constrained output, including an R-state classification (R1 to R4) and a concise explanation of the observed condition.

This design enables the model to perform guided reasoning grounded in engineering criteria, improving consistency across different inspection scenarios. Because the assessment logic is defined externally, the system remains transparent and easily adaptable, allowing domain experts to review or refine the criteria without modifying the underlying model. Overall, the rubric-based method provides a practical way to combine domain knowledge with vision-based AI, supporting reliable and interpretable bridge condition assessment. The rubric-based assessment consists of three components:

(1) **Component-specific rubric definition**: Each structural component is associated with predefined damage indicators and corresponding condition states.
(2) **Prompt construction**: Inspection images are combined with component context and rubric criteria to form structured prompts.
(3) **Vision–LLM inference**: The model generates schema-constrained outputs including classification and reasoning.

For example, a column image exhibiting spalling and exposed reinforcement may be mapped to R3 based on rubric-defined thresholds.

4. Mobile Field Deployment and Workflow

To address the limitations of traditional inspection workflows, SICA introduces a mobile-centric field deployment strategy that enables end-to-end bridge inspection directly on-site. Fig. 2 illustrates the contrast between the conventional desktop-based workflow and that of the proposed mobile-based one. In the traditional desktop-centric approach, field inspectors first capture photos during site visits and subsequently transfer the data to an office environment for processing. As shown on the left side of Fig. 2, this process introduces a separation between data collection and analysis, requiring additional steps for data transportation and upload. Such delays hinder timely decision-making, particularly in post-earthquake scenarios where rapid assessment is critical. In contrast, the proposed mobile-centric workflow, shown on the right side of Fig. 2, integrates data collection and analysis within a single platform. Using a smartphone-based interface, inspectors can capture images directly in the field, organize them by structural component, and immediately initiate the assessment process. This eliminates intermediate data handling steps and enables near real-time evaluation.

The mobile application is implemented as a web-based interface optimized for smartphone devices, enabling lightweight deployment without requiring installation through platform-specific app stores. As shown in Fig. 3, the application can be added to the device home screen and accessed similarly to a native mobile app. This approach ensures compatibility across both iOS and Android devices while maintaining ease of access for field inspectors. By providing a familiar app-like interface, the system lowers the barrier to adoption and supports efficient use in time-critical post-earthquake scenarios. Following initialization, the workflow proceeds through structured data entry and component-based inspection, ensuring consistent association between images and structural elements. Detailed steps of SICA workflow are discussed in the following paragraphs.

Following installation, the inspection workflow begins with session initialization, as illustrated in Fig. 4. The user first logs into the application, after which bridge metadata is entered, including bridge identification and location information. The interface integrates GPS-based geolocation and map visualization, allowing users to verify or adjust the bridge location directly within the app. A confirmation screen then summarizes the session details before the inspection begins, providing a clear transition between setup and data collection. This structured initialization process helps ensure data consistency and reduces the risk of associating inspection records with incorrect assets.

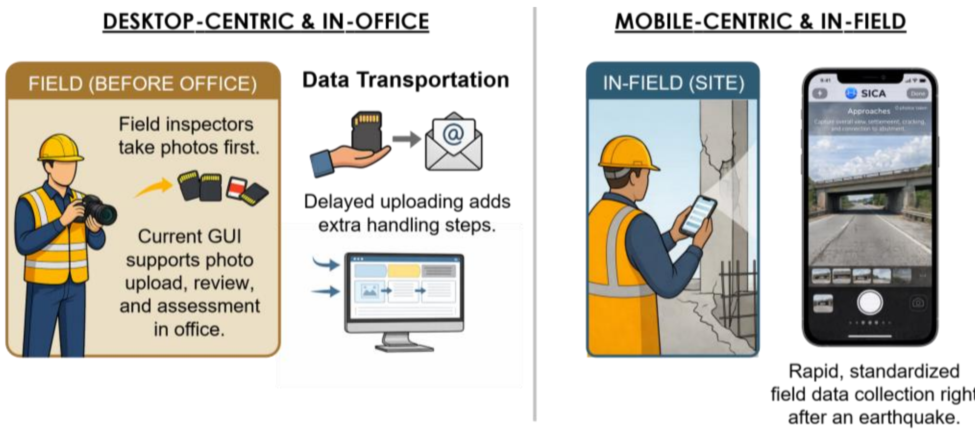

Fig. 2: Comparison of traditional desktop-based and proposed mobile in-field inspection workflows.

Fig. 3: SICA mobile web application installed on a smartphone home screen.

After session initialization, the inspection proceeds through a component-based workflow, as shown in Fig. 5. The interface organizes data collection by predefined bridge components, such as approaches, columns, joints/hinges, and abutments, ensuring that each photo is associated with the appropriate structural category. For each component, users can initiate photo collection through a dedicated entry point within the inspection page, promoting a structured and systematic inspection process.

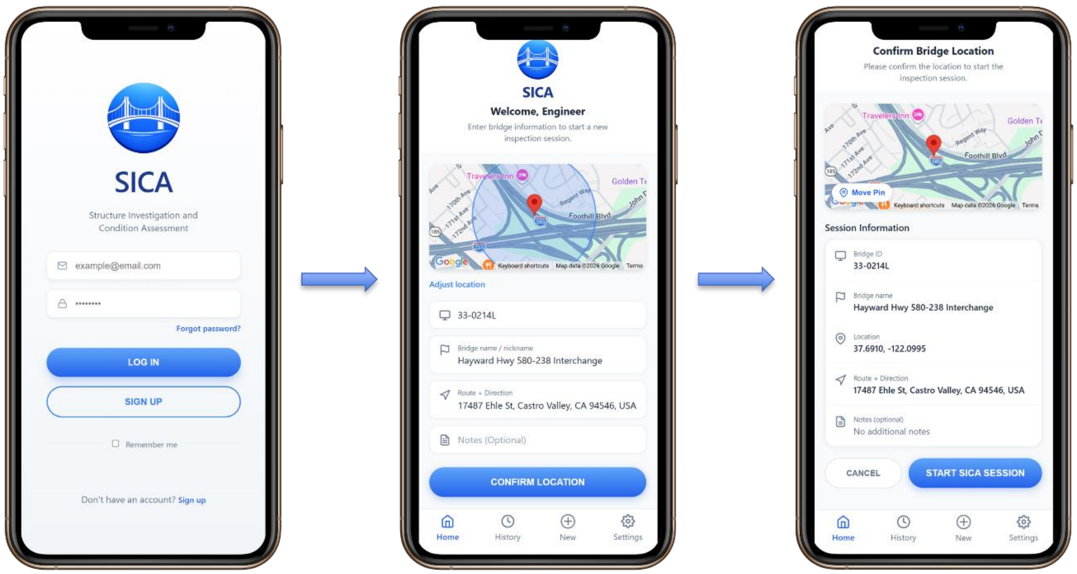

(a) Login (b) Bridge info + GPS/map (c) Confirm & start

Fig. 4: Mobile app workflow for session initialization: (a) user login, (b) bridge information entry with GPS/map integration, and (c) confirmation before starting the inspection session.

The system provides flexibility to either capture new images using the device camera or upload existing photos from the gallery, accommodating different field conditions and user preferences. During photo capture, the interface presents component-specific guidance on desired viewpoints and damage indicators, helping to ensure that collected images are both relevant and informative for subsequent assessment. In addition, the inspection interface includes lightweight progress tracking, such as the number of photos collected per component and completion status, allowing users to verify sufficient coverage before initiating analysis. This reduces the likelihood of missing critical components or incomplete documentation. By guiding users through a consistent and intuitive workflow, the system improves data quality, enhances usability in field environments, and aligns closely with established engineering inspection practices.

Following data submission, the system performs automated assessment and presents the results through a structured interface, as shown in Fig. 6. The application provides a hierarchical visualization of assessment outputs, beginning with an overall bridge-level condition classification (R1 to R4), which supports rapid decision-making such as traffic restrictions or closure. In addition to the overall assessment, the interface displays component-level results, allowing users to identify which structural elements contribute most significantly to the final decision. For each component, representative images and associated condition states are presented to highlight critical areas of concern.

The interface further supports detailed photo-level inspection, where users can explore individual images along with their corresponding R-state classifications and concise engineering rationale. This enables users to trace high-level decisions back to the underlying visual evidence, improving transparency and supporting human-in-the-loop validation. By combining multi-level summaries with detailed inspection views, the system facilitates efficient triage while maintaining clear traceability to the original data. To further support detailed review, the system provides a photo-level inspection interface, as shown in Fig. 7. Each image is assigned to an R-state classification (R1 to R4) along with a concise engineering interpretation describing the observed condition. The results are organized using a *severity-based prioritization*, where the most critical observations are presented first, enabling users to quickly identify potential safety concerns.

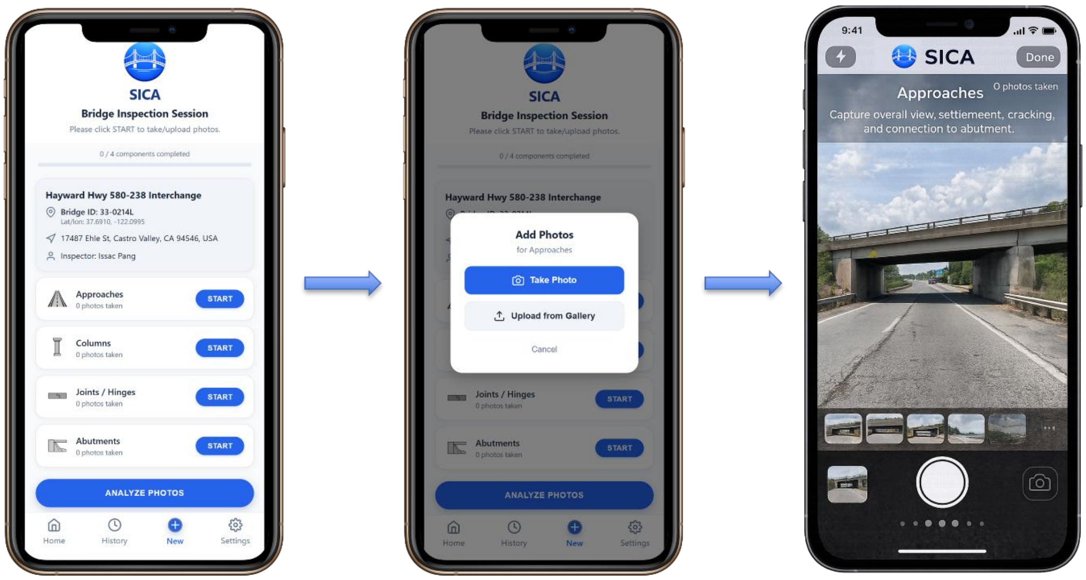

(a) Inspection page (b) Choose camera vs upload (c) Take photos

Fig. 5: Component-based inspection workflow in SICA mobile app: (a) inspection interface organized by bridge components, (b) selection of camera capture or photo upload, and (c) guided in-field image capture.

In addition to classification, the interface integrates visual evidence with explanatory reasoning, allowing users to directly associate assessment outcomes with the corresponding images. This design improves interpretability and facilitates efficient verification of automated results. By enabling users to trace component- and bridge-level decisions back to individual photos, the system supports engineering judgment, validation, and auditability within the inspection workflow.

5. Illustrative Examples of SICA Assessment

Fig. 8 shows three examples of SICA photo-level assessment using public U.S. Geological Survey images of bridge damage after earthquakes. Panels (a) and (b) are from 1994 Northridge earthquake. Panel (c) shows the Cypress Viaduct failure from 1989 Loma Prieta earthquake. These examples demonstrate how SICA converts visible damage features into clear triage decisions backed by concise engineering reasoning.

The major bridge failure presented in Fig. 8a is classified by SICA as R3 due to a major drop at the approach (bridge end) as well as apparent settlement resulting in a large discontinuity at the bridge end, making it unsafe for vehicles to drive across. An R3 state requires closure or restricted access until further inspection is conducted. The example in Fig. 8b shows a column base failure where there is extensive concrete failure, major loss of section, and the presence of extensive and exposed reinforcement indicating loss of confinement within the concrete section and a near collapse state, consistent with R4 SICA classification. Fig. 8c illustrates a multiple column failure where several columns have severe loss of concrete and exposed reinforcement. Accordingly, the piers are broken or partially collapsed, indicative of imminent structural collapse (R4 by SICA). The examples show how SICA uses visible damage states to define a transparent, action-oriented triage state system for rapid engineering review and post-event decision making.

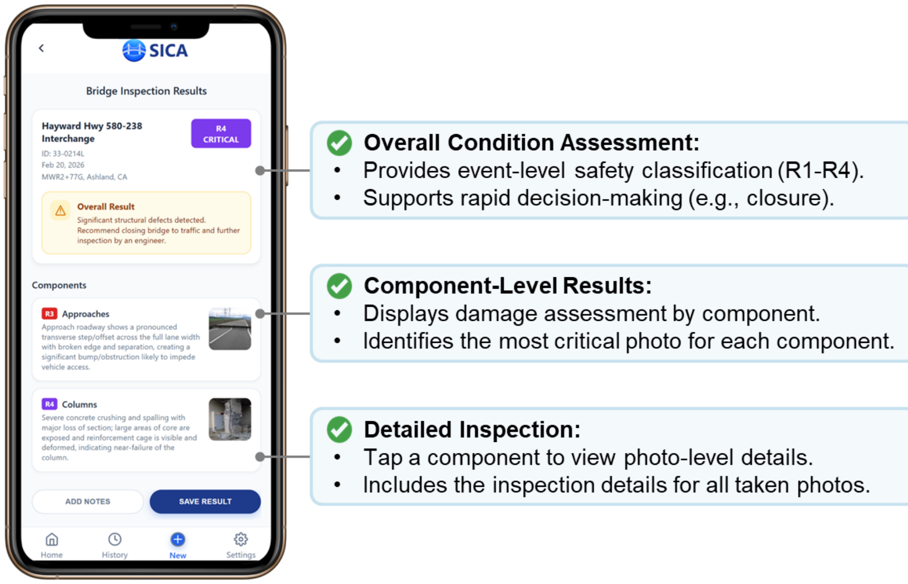

Fig. 6: Mobile app interface for assessment results: overall bridge-level condition, component-level summaries, and detailed photo-level inspection.

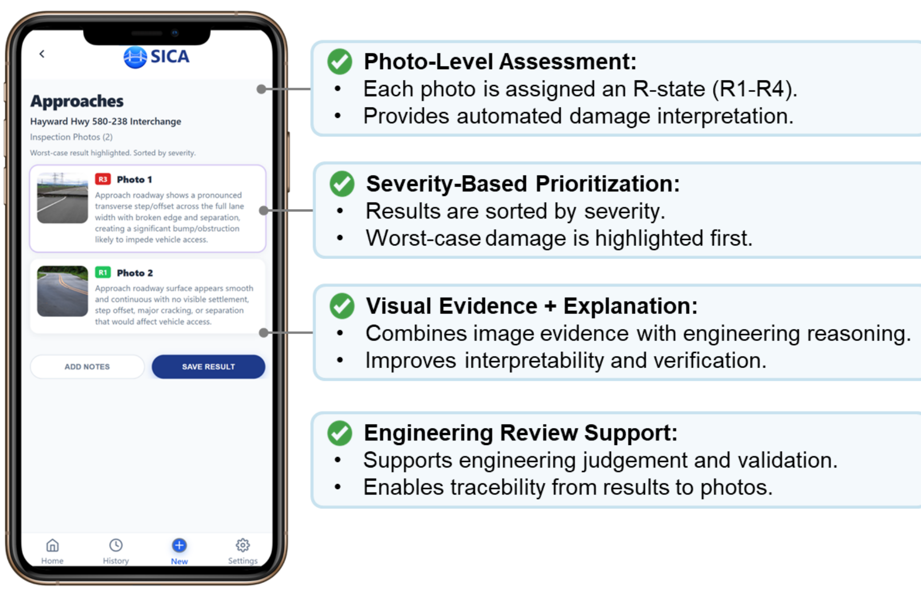

Fig. 7: Photo-level assessment interface showing severity-ranked results with R-state classification and supporting engineering reasoning.

6. Discussion

The development of SICA highlights the importance of integrating domain knowledge with AI-driven visual assessment in safety-critical applications. Rather than relying solely on data-driven models, the use of rubric-based prompting provides a mechanism to embed established engineering criteria directly into the evaluation process. This approach reflects a shift from purely predictive systems toward guided, knowledge-informed AI, where model outputs are constrained and interpreted within a structured decision framework. The results also emphasize the role of workflow design in practical deployment. While much prior work focused on improving model accuracy, the effectiveness of SICA depends equally on how data are collected, organized, and presented to the users. The component-based inspection structure and mobile interface demonstrate that consistent data acquisition and user interaction design can significantly influence the reliability and usability of AI-assisted assessment.

(a) Approach damage

major drop and separation at the bridge end with apparent settlement creating a large discontinuity that would block or severely impede vehicle access.

(b) Column base damage

Catastrophic concrete failure with major loss of section and extensive exposed/bent reinforcement, indicating loss of confinement and near-collapse condition.

(c) Multiple columns damage

Multiple columns failure with major loss of concrete section and extensive exposed reinforcement; several piers are broken/partially collapsed indicating imminent instability.

Fig. 8. Illustrative examples of SICA photo-level assessment on representative post-earthquake bridge damage images. (a) Approach damage classified as R3, indicating closure or restricted access pending further inspection. (b) Column base damage classified as R4, indicating severe instability and immediate closure. (c) Multiple-column viaduct damage classified as R4, indicating catastrophic instability and immediate closure. Source images: U.S. Geological Survey, public domain.

Another important consideration is the need for careful balance between automation and human judgment. SICA is not intended to replace engineering decision-making but to support it by providing rapid, structured summaries of visual evidence. The inclusion of concise reasoning and traceability to photo-level observations enables engineers to interpret, verify, and override automated outputs when necessary. This human-in-the-loop design is particularly important in post-earthquake scenarios, where uncertainty and variability are inherent.

The system also reveals trade-offs between standardization and flexibility. While predefined rubrics improve consistency and comparability across inspections, they may limit the ability to capture atypical or complex damage conditions. This suggests that future systems may benefit from adaptive or extensible rubric designs that evolve with additional data and expert feedback. Overall, the SICA framework illustrates that effective deployment of AI in structural engineering requires not only advances in model capability, but also careful consideration of interpretability, workflow integration, and user interaction.

Despite its advantages, the proposed framework has limitations. The performance of the vision-LLM system depends on image quality, coverage, and the completeness of rubric definitions, as mentioned above.

In addition, the current study does not include large-scale quantitative validation against labeled inspection datasets, which remains an important direction for future work.

7. Conclusion

This paper presented SICA, a framework for rubric-based vision-LLM assessment of bridge inspection images, integrated with a mobile-centric workflow for in-field deployment. By combining structured engineering knowledge with vision-based AI, the system enables interpretable and standardized condition evaluation from inspection photos. The proposed mobile interface supports end-to-end inspection directly in the field, including data collection, automated assessment, and multi-level result visualization. This integration reduces data handling overhead, improves consistency in inspection workflows, and enables more timely decision support in post-earthquake scenarios. Overall, SICA demonstrates the potential of combining AI-driven visual assessment with structured engineering workflows to support rapid and reliable bridge condition screening. The framework provides a foundation for future development toward scalable, data-informed, and field-deployable infrastructure assessment systems. Specifically, future work will focus on large-scale validation, integration with sensor-based structural health monitoring data, and deployment in real-world inspection workflows in collaboration with infrastructure agencies and industry partners.

8. Reference

[1] California Department of Transportation, “First responder bridge assessment guide: Structure maintenance & investigations (SM&I),” Div. of Maintenance, Office of Emergency Management and Infrastructure Protection, 2016.

[2] K.M. Mosalam and Y. Gao, “Artificial intelligence in vision-based structural health monitoring,” Cham, Switzerland: Springer, 2024.

[3] P.J. Vardanega, T. Tryfonas, G. Gavriel, D.T. Nepomuceno, M. Pregnolato, and J. Bennetts, “A review of recent research on visual inspection processes for bridges and the potential uses of AI,” in Bridge Maintenance, Safety, Management, Digitalization and Sustainability, CRC Press, 2024.

[4] T. Panigati, M. Zini, D. Striccoli, P. F. Giordano, D. Tonelli, M. P. Limongelli, and D. Zonta, “Drone-based bridge inspections: Current practices and future directions,” Automation in Construction, vol. 173, p. 106101, 2025.

[5] K. Luo, X. Kong, J. Zhang, J. Hu, J. Li, and H. Tang, “Computer vision-based bridge inspection and monitoring: A review,” Sensors, vol. 23, no. 18, p. 7863, Sep. 2023.

[6] Z. Li, X. Wu, H. Du, F. Liu, H. Nghiem, and G. Shi, “A Survey of State of the Art Large Vision Language Models: Alignment, Benchmark, Evaluations and Challenges,” in 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2025

[7] S. Danish, A. Sadeghi-Niaraki, S.U. Khan, L.M. Dang, L. Tightiz, and H. Moon, “A comprehensive survey of Vision–Language Models: Pretrained models, fine-tuning, prompt engineering, adapters, and benchmark datasets,” Information Fusion, vol. 126, p. 103623, 2026

[8] A. Sharshar, L.U. Khan, W. Ullah, and M. Guizani, “Vision-language models for edge networks: A comprehensive survey,” IEEE Internet of Things Journal, vol. 12, no. 16, p. 32701, Aug. 2025.