


Article

Effectiveness of Generative AI for Post-Earthquake Damage Assessment

João M. C. Estêvão ^{1,2} ¹ ISE—University of Algarve (UAlg), Campus da Penha, 8005-139 Faro, Portugal; jestedvao@ualg.pt² CIMA—Centre for Marine and Environmental Research, ARNET—Infrastructure Network in Aquatic Research, UAlg, Campus de Gambelas, 8005-139 Faro, Portugal

Abstract: After an earthquake, rapid assessment of building damage is crucial for emergency response, reconstruction planning, and public safety. This study evaluates the performance of various Generative Artificial Intelligence (GAI) models in analyzing post-earthquake images to classify structural damage according to the EMS-98 scale, ranging from minor damage to total destruction. Correct classification rates for masonry buildings varied from 28.6% to 64.3%, with mean damage grade errors between 0.50 and 0.79, while for reinforced concrete buildings, rates ranged from 37.5% to 75.0%, with errors between 0.50 and 0.88. Fine-tuning these models could substantially improve accuracy. The practical implications are significant: integrating accurate GAI models into disaster response protocols can drastically reduce the time and resources required for damage assessment compared to traditional methods. This acceleration enables emergency services to make faster, data-driven decisions, optimize resource allocation, and potentially save lives. Furthermore, the widespread adoption of GAI models can enhance resilience planning by providing valuable data for future infrastructure improvements. The results of this work demonstrate the promise of GAI models for rapid, automated, and precise damage evaluation, underscoring their potential as invaluable tools for engineers, policymakers, and emergency responders in post-earthquake scenarios.

Keywords: post-earthquake damage assessment; generative artificial intelligence; damage classification; EMS-98 scale



Citation: Estêvão, J.M.C. Effectiveness of Generative AI for Post-Earthquake Damage Assessment. *Buildings* **2024**, *14*, 3255. <https://doi.org/10.3390/buildings14103255>

Academic Editor: Mizan Ahmed

Received: 12 September 2024

Revised: 9 October 2024

Accepted: 10 October 2024

Published: 14 October 2024



Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In regions with a history of destructive earthquakes, it is vital for communities to be prepared to respond effectively to such natural disasters. Improving the ability to quickly and safely assess the seismic vulnerability of affected buildings, based on the damage they have sustained, plays a crucial role in ensuring a swift and effective response.

In large-scale disasters, such as the earthquake that struck Turkey in 2023, thousands of individuals may be required to assess the degree of damage to buildings affected by an earthquake [1]. This assessment is essential for ensuring that buildings can be safely used and must be conducted as quickly as possible, often under emergency conditions. However, post-earthquake damage assessment is a complex task, which is particularly challenging when relying solely on visual inspection, given the various types of damage that may occur in buildings affected by seismic vibrations [2–4].

In this context, it is crucial to explore the potential use of artificial intelligence (AI) algorithms for assessing building damage following a devastating earthquake. This approach can be particularly valuable when there is a limited number of available human resources to perform the task. AI (refer to the list of abbreviations in Abbreviations) can assist in identifying severe damage or collapses, which may involve potential victims, as well as in evaluating less critical damage that could still affect the safe occupancy of buildings. Addressing these issues promptly is essential to managing the number of displaced individuals, who also require timely assistance.

There is a wide variety of AI algorithms, which can be categorized into several subcategories (Figure 1). These AI algorithms possess different capabilities, making it essential to understand some fundamental concepts before selecting one to address a specific problem. For example, some AI algorithms are particularly effective in handling classification tasks, others are better suited for addressing regression problems, and some are optimized for solving complex optimization challenges [5,6].

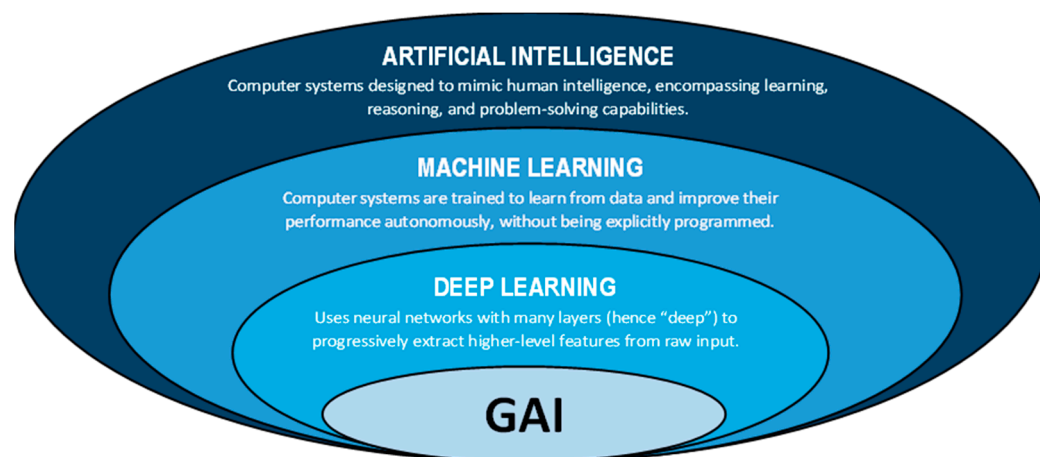


Figure 1. Diagram of the different subgroups of methods in the field of AI [6,7].

The use of AI in the context of earthquake engineering is not a new concept. For example, in 1998, Estêvão was one of the first researchers in Portugal to propose the application of AI in seismic risk assessment for buildings. Specifically, in his MSc thesis [8], he presented the idea and used a combination of a fuzzy expert system with a deep learning (DL) system based on multi-layered feedforward neural networks (NN), utilizing an Error Backpropagation Learning Algorithm. That generation of researchers was significantly influenced by the success of the AI system known as DeepBlue, which, for the first time in human history, was able to defeat the world chess champion, Kasparov, in 1997 [9]. It is important to highlight that, at that time, the wide array of AI libraries for various programming languages (such as Python or C#) that we have today did not exist. Therefore, these algorithms had to be studied in great detail before they could be implemented in software, as was the case with the NEUNET software [8]. However, this line of research often remained stagnant for many years. For example, Estêvão’s study [8] was only partially published in a journal almost 20 years after it was initially presented, now including results from nonlinear analyses used to train the neural networks (NN) [10]. This approach was later extended to the context of schools [11], following the PERSISTAH project [12]. The temporal gap in the adoption of AI for various applications was largely due to the evolution of AI itself and advancements in computational systems. In the past, there was a significant limitation in the computational power required to train deep learning (DL) systems. Additionally, a substantial portion of the scientific community was highly skeptical about the use of AI, which hindered the formation of collaborations, the acquisition of funding, and made it challenging to get AI-related research in earthquake engineering published in scientific journals. This resistance was also evident in other scientific fields, such as computer vision. A pertinent example is a paper submitted to a conference by Yann LeCun and his collaborators, which was rejected because the scientific community at the time did not believe that deep learning systems could learn to classify images solely from examples [13]. It is important to highlight that Yann LeCun was a pioneer in using Convolutional Neural Networks (CNN) for image recognition [14], and today he is a well-known and highly awarded researcher. In essence, there were two prevailing perspectives: on one side were the researchers who had an intuition about the potential of DL, and on the other side was the majority of researchers who were highly distrustful of these capabilities,

primarily because they viewed these systems as “black boxes” where humans could only control the training process of the AI system, not what it learned.

A significant turning point in the perception of DL capabilities occurred in 2012 with the advent of the widely recognized AlexNet [13,15], which demonstrated its effectiveness in accurately classifying various types of images. This shift in perception was further reinforced in 2016 when AlphaGo defeated a world champion human player in the game of Go, followed by AlphaZero’s remarkable performance in chess [9]. However, DL systems require substantial amounts of data to ensure the proper training of neural networks; otherwise, issues like overfitting may arise [10]. In the domain of computer vision, strategies must also be developed to combat overfitting, especially when there is an insufficient number of images available, such as employing distinct forms of data augmentation [13].

With the emergence of the transformer architecture in 2017 [16], followed by its application in natural language processing (NLP) [17], generative artificial intelligence (GAI) gained significant visibility within the scientific community. This visibility extended to the general public in November 2022 with the release of ChatGPT, showcasing its potential across various domains [6]. Today, the conversational abilities and cognitive performance exhibited by the most recent large language models (LLMs) are already evident [18], particularly in the chatbots currently in use.

In the context of post-earthquake damage assessment, several methods for the rapid and automatic evaluation of building damage were recently tested. The first attempt to use remote sensing imagery for earthquake damage mapping, still preliminary in nature, occurred during the 1906 San Francisco earthquake in California, USA, where aerial photographs of the affected areas were obtained. The first use of satellite images likely dates back to 1972, concerning the 1964 Alaska earthquake [19]. Currently, the methods typically proposed rely on satellite images, aerial views, images captured by unmanned aerial vehicles (UAVs) at the building level, or combinations of these image sources [19–23]. Many of these methods employ various types of AI algorithms, such as Support Vector Machines (SVM) and Random Forest (RF), which are ML algorithms, or CNNs, which are DL algorithms [19,24–26].

RF is an algorithm typically used to solve regression or classification problems. In the context of post-earthquake scenarios, these algorithms can be trained to classify the extent of building damage into categories such as “standing” and “collapsed” over large areas using imagery [27].

SVM is a supervised learning algorithm primarily used in classification problems, allowing the assignment of a damage grade based on the distance of a result to a hyperplane that divides the data into classes [28].

The use of CNNs is, very likely, the most widely adopted strategy for assessing building damage in post-earthquake scenarios, with numerous recent studies published on the topic [29–43]. Various network architectures based on AlexNet [31], previously recognized as a significant milestone in computer vision, were developed, alongside many other CNN architectures [30]. An important conclusion is that CNN-based networks encounter several challenges, particularly related to the loss of spatial information and the fact that training data may not fully represent the problem domain [30].

Generative Adversarial Networks (GANs) are relatively recent models of GAI that are beginning to be employed in earthquake engineering [44], particularly for post-earthquake damage assessment in buildings [45]. In this context, a pertinent question arises: Can GAI, especially considering the current availability of multimodal systems with models that possess an enormous number of parameters and were trained on supercomputers using vast amounts of data from the internet, effectively contribute to solving the problem of damage assessment in post-earthquake scenarios? Addressing this question, the present work conducts a preliminary study aimed at providing an initial answer to this inquiry.

2. GAI Models

Given that the approach proposed in this work involves GAI concepts that are not yet widely understood by all stakeholders in post-earthquake damage assessment, and because information on the subject is still fragmented, this section was created to provide the necessary clarifications within the context of the research conducted, which involves both text and image data.

For a GAI system to effectively assess the damage level of a building in a post-earthquake scenario based on a single image, it is likely desirable to utilize a Foundation Model (FM) [46,47] with multimodal capabilities (text–image–text). Access to such models can be facilitated through a chatbot, which provides a user-friendly interface, or via an Application Programming Interface (API), which is more suited for software developers.

Since the introduction of the first version of ChatGPT, a chatbot powered by a Generative Pretrained Transformer (GPT)-based LLM, the potential impact of such models on society has become evident, encompassing both positive and negative aspects [48]. To enable these chatbots with multimodal capabilities [6], multimodal models can be employed, such as a text-to-image diffusion model like DALL-E 3 [49], or a multimodal LLM with text-to-text and image-to-text capabilities like GPT-4 [50]. Alternatively, a combination of unimodal models may be used (e.g., an LLM like GPT-3.5 [51], initially introduced in ChatGPT on 22 November 2022 [52], paired with an image generation model), or hybrid approaches may be adopted that integrate certain multimodal components with unimodal ones, working together synergistically (Figure 2).



Figure 2. GAI models (examples include text-to-text, text-to-image, image-to-text, and image-to-image). The image of the damaged building that is shown in Figure 2 was generated by the DALL-E 3 model [53], to illustrate the current capabilities of GAI.

Given the multiplicity of an FM and the rapid emergence of new models, it is crucial to understand their differences. One significant characteristic is the model size, which has increased with each generation, necessitating larger training datasets to prevent overfitting. For example, the GPT-1 model was developed with 117 million parameters and trained on 4.5 gigabytes of text data, while GPT-2 expanded to 1.5 billion parameters and was trained on 40 gigabytes of text data. GPT-3 marked a substantial leap with 175 billion parameters and was trained on 570 gigabytes of text [6,51]. As for GPT-4, its exact size is debated, with

some suggesting it consists of trillions of parameters [6,51], though this number has not been confirmed by OpenAI [50].

The performance of these models also varies significantly. For instance, GPT-3.5's performance is notably inferior to GPT-4 across various benchmarks [50], highlighting the need for careful model selection when using AI to assist with specific tasks.

On the other hand, GAI models can be categorized into two major groups: open-source models (or open-weight models) and closed-source models [54]. Open-source models are free and available for anyone to use, modify, and distribute. In the case of open-weight models, only the model's parameters (the neural network weights) are modifiable, not the underlying code. Typically, closed-source models are larger and offer superior performance, but they limit customization and provide less control over the training data. However, there are now large open-source models with performance comparable to those of closed-source models [47]. The choice between an open-source (or open-weight) model and a closed-source model depends on the specific needs of its use, including budget, security requirements, and the need for customization and control over data and the model.

2.1. Transformer Architecture

The emergence of the transformer architecture [16] and GPT [17] has been pivotal in the development of LLMs. Therefore, it is crucial to have a fundamental theoretical understanding of these topics, particularly regarding their evolution, to better anticipate potential application domains [6].

2.1.1. Tokenization

Tokenization is the process aimed at creating a digital representation of a real-world object. In the context of LLMs, text is divided into tokens, which can be words, sub-words, or individual characters. These tokens are mapped to specific codes (numerical identifiers) [55], as illustrated in Figure 3. The same principle can also be applied to images, where images are similarly tokenized to facilitate processing by AI models [56].

Generative Pretrained Transformers (GPT)
[5648, 1413, 5075, 36822, 81632, 320, 38, 2898, 8]

Figure 3. Example of text tokenization in GPT-4 [55]. The text associated with each color corresponds to a token.

2.1.2. Embeddings

Each token is converted into a numerical vector known as an embedding, of a specific dimension (dimension N_E), as depicted in Figure 4. This vector aims to represent the meaning of the word (or token) in a multidimensional vector space, capturing both semantic and syntactic relationships with other tokens [57]. Accordingly, the model has a predefined vocabulary of tokens with a certain size (N_V), and an embedding matrix is created with dimensions equal to $N_E \times N_V$ (each column represents the vector representation of a token from the vocabulary in the embedding space) [58].




Figure 4. Example of embedding creation.

The distance between two vectors measures their relatedness [59,60]. Small distances indicate high relatedness, while large distances suggest low relatedness. To better understand this concept, consider the embeddings represented in the two-dimensional space

shown in Figure 5. The “Royalty” cluster consists of the words King, Queen, and Prince, which are located close to each other because they are semantically related, representing terms associated with monarchy and nobility. Their coordinates reflect their strong semantic connections, forming a tight cluster in the vector space. The “Tropical Fruits” cluster is composed of the words Pineapple and Papaya, which are located in another region of the space, reflecting their semantic relationship as tropical fruits. This cluster is distant from the Royalty cluster, highlighting the distinct context and meaning of these words. The “Geographical” cluster contains only the word Tropics, which is associated with geographical and environmental concepts, particularly relating to regions where fruits like pineapples and papayas grow. Therefore, it is placed closer to the Tropical Fruits cluster than to the Royalty cluster but still distinct to indicate its unique context.

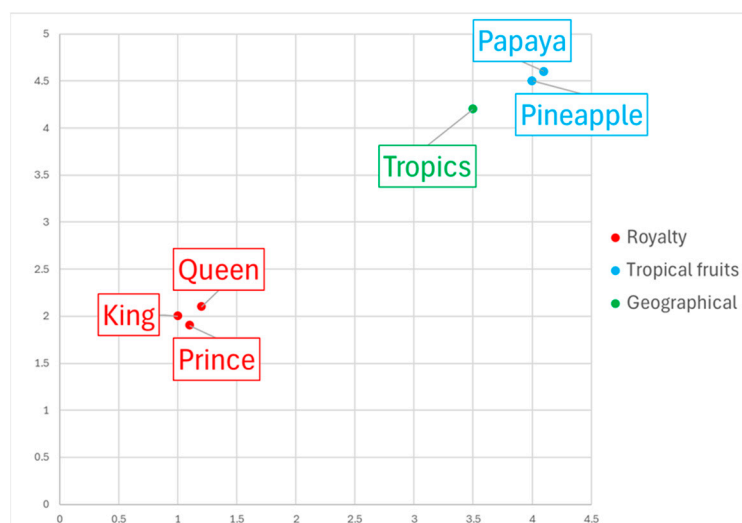


Figure 5. Illustration of a possible spatial representation of embeddings (in a 2-dimensional space).

2.1.3. Attention Mechanism

There are different types of attention mechanisms, such as self-attention, which occurs within the same sequence (e.g., between words in a sentence), and cross-attention, which occurs between different sequences (e.g., between a question and a context in question-answering systems) [61].

Typically, attention mechanisms involve three distinct steps in each attention layer: calculating the attention score, applying the softmax function to obtain probabilities, and generating the context vector.

The attention score is a concept in DL models that combines proximity in the embedding space and relevance in the context. For each token, three vectors are created [16]: Query (Q), Key (K), and Value (V), using projection matrices W_Q , W_K , and W_V (weight matrices whose dimensions correspond to the embedding space dimensions and the projection space dimensions, which typically match the embedding space dimensions) [61].

The Q vector represents the query for which the model is attempting to find relevant information. It is derived from the token being processed and seeks to identify which parts of the input should be focused on. The K vector represents the “key” that is compared with the query to measure compatibility or relevance. Each token in the input sequence has its own key vector. The V vector contains the information that will be extracted and combined based on the relevance calculated between the query and the keys. It represents the information associated with each token [61].

For each token i and each layer of the transformer model, we have the following:

$$Q_i = \text{embedding_token}_i \cdot W_Q, \quad (1)$$

$$K_i = \text{embedding_token}_i \cdot W_K, \quad (2)$$

$$V_i = \text{embedding_token}_i \cdot W_V. \quad (3)$$

Many practical implementations utilize multi-head attention mechanisms [16], where multiple sets of Q , K , and V vectors are used in parallel.

In each layer of the transformer model, when evaluating a token, the relevance (or “attention”) between that token and all other tokens in the prompt is calculated. Therefore, an attention score matrix is generated (with dimensions $N_T \times N_T$, where N_T is the number of tokens in the prompt). Each entry of the matrix $score_{ij}$ represents the attention score between token i and token j in the prompt and is computed as the following:

$$score_{ij} = Q_i \cdot K_j^T. \quad (4)$$

Subsequently, the attention scores are normalized using the softmax function, converting them into probabilities. These probabilities sum to 100% and indicate the relative importance of each token in the context [58].

The context vector for a token i is then calculated as a weighted combination of the V vectors of all tokens, using the attention probabilities as weights, given by:

$$\text{context_vector}_i = \sum_{j=1}^{N_T} \text{attention_weights}_{ij} \cdot V_j \quad (5)$$

These vectors serve as the input to the subsequent layer. The process is repeated until the output of the last attention layer is obtained [58].

It is important to highlight that these models can only process a certain number of tokens, a characteristic referred to as the context window [62], which is a crucial feature of a GAI model.

2.1.4. Next Token Prediction

The process of predicting the next token involves several steps. The context vector, derived from the attention mechanism, is processed through additional layers of the NN, such as feed-forward neural network layers. The resulting state then passes through an output layer, typically a linear transformation followed by a softmax function. This generates a probability distribution over all tokens in the model’s vocabulary. The token with the highest probability is generally selected as the next in the sequence, though techniques like sampling can be employed to introduce diversity. This process occurs iteratively, with each new predicted token being added to the context for the next prediction [47,58,61] (Figure 6).




Figure 6. Illustration of the next token prediction process in a typical architecture [47].

2.1.5. Temperature and Top-P

Temperature is a parameter that controls the randomness of predictions made by an LLM. It adjusts the probability distribution associated with the prediction of the next token. A lower temperature (<1) reduces randomness, leading the model to select words with higher probabilities more conservatively. This property results in more predictable and deterministic text. If the temperature is set to zero, the next token generated will be the one with the highest probability. Conversely, a higher temperature (>1) increases randomness, allowing the model to choose fewer probable words more frequently. This characteristic can produce more varied and creative text but may also introduce incoherence [58].

In addition to temperature, another parameter that can be adjusted in a model (especially when using an API) is top-P (nucleus sampling). Instead of considering all possible words, the model selects only the words whose cumulative probability reaches a cut-off value P [63].

Low values of temperature and top-P result in safer and more coherent texts, which are ideal for applications where precision and clarity are crucial, such as document summaries or formal responses. Higher values of temperature and top-P allow for more creative and varied text generation, making them better suited for tasks like story writing or brainstorming.

2.2. Prompt Engineering

To develop effective strategies for utilizing GAI models in socially beneficial activities, it is essential to first have a theoretical understanding of the subject. This understanding involves leveraging human creativity and possessing knowledge of prompt engineering (Figure 7).

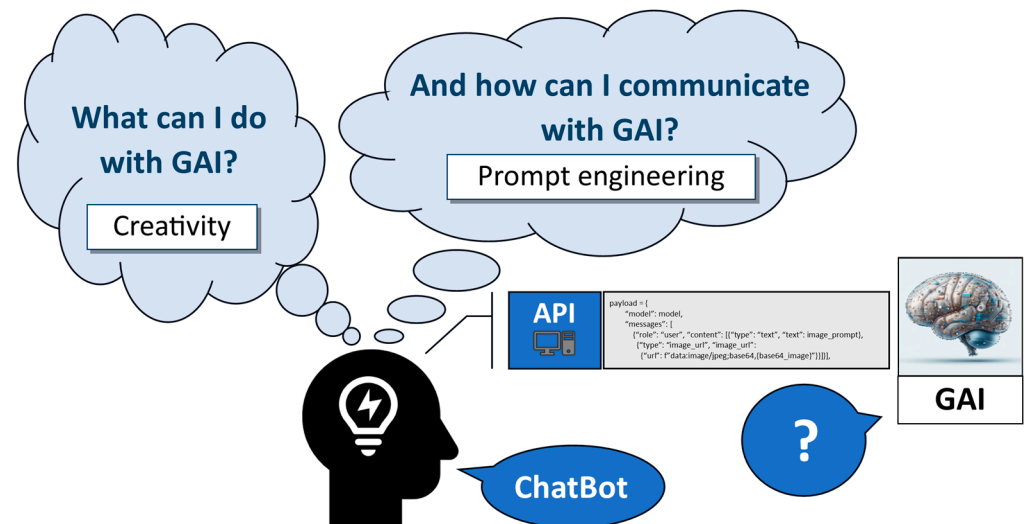


Figure 7. Diagram illustrating possible interaction between a human and a GAI model.

A prompt is a set of instructions given to a GAI model, such as a Large Language Model (LLM), to generate a desired output. The prompts can significantly influence the generated results, including their accuracy. In this context, prompt engineering is becoming increasingly important [64], particularly in efforts to minimize the occurrence of “hallucinations” in LLMs.

In the context of LLMs, a “hallucination” refers to the generation of incorrect or unintended responses [65]. These hallucinations may present in various forms: (1) internal inconsistency within the response, such as the LLM generating contradictory statements; (2) contradiction of the prompt, where the response diverges from the intended meaning of the question (e.g., a positive inquiry receives a negative reply); (3) factual inaccuracies; or (4) contextually misplaced information, where the correct data are presented but does not adhere to the specific instructions provided to the LLM.

Several factors contribute to hallucinations in LLMs. The most evident is the quality of the training data. If the LLM lacks data on a particular topic, it is more likely to generate a hallucination-producing text that is grammatically correct but factually inaccurate. Thus, the more specialized the topic, particularly if it is not well represented in freely accessible online sources, the higher the likelihood of hallucination. The type of LLM also plays a role; for instance, the hallucination rate is higher in GPT-3.5 compared to GPT-4 [66]. Another significant source of hallucinations relates to the quality of the prompts used. A well-structured prompt can guide the LLM toward generating accurate results, whereas

a poorly constructed prompt might lead to hallucination [64]. This aspect is the part of hallucination control that users can most easily manage.

In this context, the need for the development of prompt engineering emerges, where human capabilities such as creativity and intuition play a crucial role, alongside interdisciplinary knowledge [67].

A good prompt should be concise (with brief and clear instructions), logical (structured and coherent instructions), explicit (clear directives on how to present results), adaptive (balancing creativity and specificity), and reflective (prompt refinement should result from multiple tests that assess their effectiveness in known domains, particularly regarding the accuracy, coherence, and utility of responses within the given instructions) [67].

The so-called zero-shot prompt is the most basic type of instruction, where no specific context is provided on the subject being queried, yet the LLM is expected to generate a suitable response using its general knowledge (e.g., “Describe what an earthquake is.”).

One way to enhance LLM outcomes is by utilizing multiple-shot prompts. In this type of prompt, several instructions are provided to give the LLM the necessary context to recognize patterns by presenting a logical sequence of examples and instructions, thereby reducing the likelihood of hallucination [68]. One-shot prompts are a particular case where only one example is given for the LLM to correctly replicate that pattern in another task. This type of prompt requires some subject matter knowledge to ensure that the example presented to the LLM is accurate.

More sophisticated strategies are also emerging, such as the use of “Tree of Thoughts” [69], where instructions are provided in a manner that encourages the LLM to create different lines of reasoning (which can involve various entities, such as virtual personas). These strategies can lead to a more consensual and potentially more accurate solution.

Thus, the user’s creativity in crafting prompts, combined with their knowledge of prompt engineering, seems to be factors that can influence the reliability of results obtained using multimodal generative AI models, making it essential to test and evaluate these approaches in the context of building damage assessment.

2.3. Fine-Tuning

Typically, multimodal AI models are built upon an FM that has been trained on large datasets, often sourced from the internet. These models then undergo a fine-tuning process, often leveraging Reinforcement Learning (RL), which enhances their initial capabilities [50]. There are various strategies for implementing this fine-tuning [70].

Focusing on the current capabilities of multimodal models to extract text, objects, and general concepts from images, it becomes evident that fine-tuning the model for the specific application domain under consideration could significantly improve the accuracy of the results. However, this task presents challenges: it is computationally intensive, potentially leading to high costs, and currently, there are few multimodal models that allow for further fine-tuning. For cost-effectiveness, one might need to select a smaller, open-source (or open-weight) multimodal model. Therefore, exploring alternative strategies to enhance model performance is crucial.

2.4. Retrieval-Augmented Generation (RAG)

RAG is a technique that combines the ability to retrieve relevant information (Retrieval), interpret and augment it (Augmented), and generate well-formulated and complete responses (Generation) [71]. This technique closely mimics how a human would respond to a specific question—by reading relevant material, particularly when they do not have full knowledge on the subject.

RAG can be successfully adapted to specific problems [72], particularly in enhancing the accuracy of image interpretation, such as in radiological images [73]. Therefore, it is crucial to assess its effectiveness in post-earthquake damage evaluation using images.

3. Post-Earthquake Damage Assessment

One of the most critical tasks in the context of civil protection following an earthquake is the assessment of structural damage to buildings. This systematic assessment is vital for guiding post-earthquake response actions, such as determining the need for resident relocation, planning emergency interventions, and supporting reconstruction efforts in affected areas. Additionally, it provides valuable data to improve construction practices and strategies for mitigating seismic risks [74].

To streamline the execution of this critical task, it is common for civil protection organizations in countries located in earthquake-prone regions to develop forms for damage assessment. In Europe, it is common to use the damage grades from the European Macroseismic Scale 1998 (EMS-98) [75] to support this assessment.

EMS-98 is an important tool for assessing the impacts of earthquakes. This scale classifies building damage into six levels, ranging from D0 to D5, describing damage severity progressively: D0 (No Damage)—no visible damage after a perceptible earthquake; D1 (Slight Damage)—small cracks in walls and superficial damage to non-structural elements like plaster and ceilings; D2 (Moderate Damage)—more pronounced cracks in walls, the detachment of plaster, and significant damage to non-structural elements such as chimneys; D3 (Substantial Damage)—structural damage, including cracks in beams and columns, as well as extensive damage to non-structural components; D4 (Severe Damage)—severe structural damage, potentially leading to a partial building collapse and generally requiring immediate evacuation; and D5 (Destruction)—a total or partial collapse of the building, rendering the structure completely uninhabitable.

However, conducting a rigorous assessment of a building's damage level can be a complex and time-consuming task, potentially requiring in situ testing [3]. Additionally, evaluators must be familiar with typical damage patterns across different construction systems, which can vary significantly. For instance, the damage patterns in traditional buildings, such as adobe structures [4], are distinctly different from those commonly observed in reinforced concrete (RC) structures [2].

Therefore, this type of technical task demands qualified personnel, who may be insufficient in number in the context of a widespread earthquake disaster. In this scenario, it is crucial to explore alternatives that reduce the need for human resources, making it important to test the use of AI.

4. Tested Approach

The approach proposed and partially tested in this study involves assessing post-earthquake damage by utilizing images of buildings captured on-site (e.g., using drones) and employing multimodal GAI alongside RAG or fine-tuning techniques to enhance the accuracy of the results.

4.1. Dataset and Data Preparation

Only photographs of the buildings presented in the EMS-98 document were used in this work. There were two primary reasons for selecting the images from the EMS-98 document to test the capabilities of GAI in identifying damage to buildings affected by earthquakes. First, the document is openly accessible, eliminating any concerns related to confidentiality. Second, the images contained within it are exemplary, setting the standard for this type of damage assessment. If other types of images, particularly those captured from recent earthquakes, had been used, there would always be a question regarding the accuracy of human interpretation in damage assessment, which would then be compared to the models.

4.2. Tested GAI Models

The performance of models with varying numbers of parameters and attention window sizes (with temperature = 1 and top-P = 0.95) was evaluated, specifically OpenAI's GPT-4o [76] (a large model that currently leads in many benchmarks, featuring a context

window of 128,000 tokens), Google’s Gemini 1.5 Pro [77] (a large model with an enormous context window of approximately 2 million tokens), OpenAI’s GPT-4o mini [78] (equivalent to GPT-4o but smaller, accessed through the computer code listed in Appendix A), and Google’s Gemini 1.5 Flash [77] (with a context window of about 1 million tokens).

4.3. Evaluation Metrics and Methods

The EMS-98 scale was utilized in this study to test the current capabilities of multi-modal GAI models in assessing damage to masonry and reinforced concrete (RC) buildings in post-earthquake scenarios.

The assessment of the type of structural system present in the buildings identified in the photographs from the EMS-98 document, as well as the corresponding classification of the degree of damage, was obtained using different GAI models and the prompt provided in Appendix B. The models’ performances were compared using accuracy rates and mean error margins for damage classifications.

4.4. Results

The results obtained are presented in Table 1 and Figure 8 for the images in the EMS-98 document related to masonry buildings and in Table 2 and Figure 9 for the images related to RC buildings.

Table 1. Performance evaluation of GAI models for building earthquake damage assessment using EMS-98 imagery—masonry buildings.

N.	Figure	EMS-98 Document	GPT-4o	Gemini 1.5 Pro	GPT-4o Mini	Gemini 1.5 Flash
1	5-1	D3	Masonry (85%) D3 (80%) ¹	Masonry (95%) D3 (80%)	Masonry (85%) D3 (80%)	Masonry (100%) D2 (90%)
2	5-2	D4	Masonry (90%) D2 (80%)	Masonry (95%) D2 (80%)	Masonry (85%) D3 (80%)	Masonry (100%) D3 (80%)
3	5-3	D4	Masonry (95%) D2 (90%)	Masonry (95%) D3 (80%)	Masonry (85%) D3 (80%)	Masonry (100%) D3 (100%)
4	5-4	D4	Masonry (90%) D4 (85%)	Masonry (95%) D5 (95%)	Masonry (90%) D3 (80%)	Masonry (100%) D4 (90%)
5	5-5	D5	Masonry (90%) D4 (85%)	Masonry (95%) D4 (85%)	Masonry (85%) D4 (90%)	Masonry (100%) D5 (90%)
6	5-6	D2	Masonry (90%) D2 (80%)	Masonry (95%) D1 (70%)	Masonry (85%) D2 (75%)	Masonry (100%) D2 (80%)
7	5-7	D3	Masonry (90%) D3 (80%)	Masonry (100%) D3 (80%)	Masonry (85%) D3 (80%)	Masonry (100%) D2 (90%)
8	5-8	D4	Masonry (95%) D4 (90%)	Masonry (95%) D4 (90%)	Masonry (90%) D3 (85%)	Masonry (100%) D4 (90%)
9	5-9	D2	Masonry (90%) D3 (85%)	Masonry (95%) D1 (75%)	Masonry (85%) D2 (80%)	Masonry (100%) D2 (100%)
10	5-10	D2	Masonry (90%) D2 (80%)	Masonry (95%) D1 (75%)	Masonry (90%) D2 (85%)	Masonry (100%) D1 (90%)
11	5-11	D2	Masonry (90%) D2 (85%)	Masonry (90%) D1 (80%)	Masonry (85%) D3 (80%)	Masonry (100%) D2 (90%)
12	5-12	D2	Masonry (90%) D2 (85%)	Masonry (95%) D1 (80%)	Masonry (85%) D2 (80%)	Masonry (100%) D2 (80%)
13	5-13	D3	Masonry (90%) D3 (80%)	Masonry (95%) D3 (85%)	Masonry (85%) D2 (75%)	Masonry (100%) D2 (100%)
14	5-14	D4	Masonry (90%) D3 (85%)	Masonry (95%) D3 (85%)	Masonry (85%) D3 (80%)	Masonry (100%) D3 (90%)

¹ In brackets is the degree of confidence that was presented by the model.

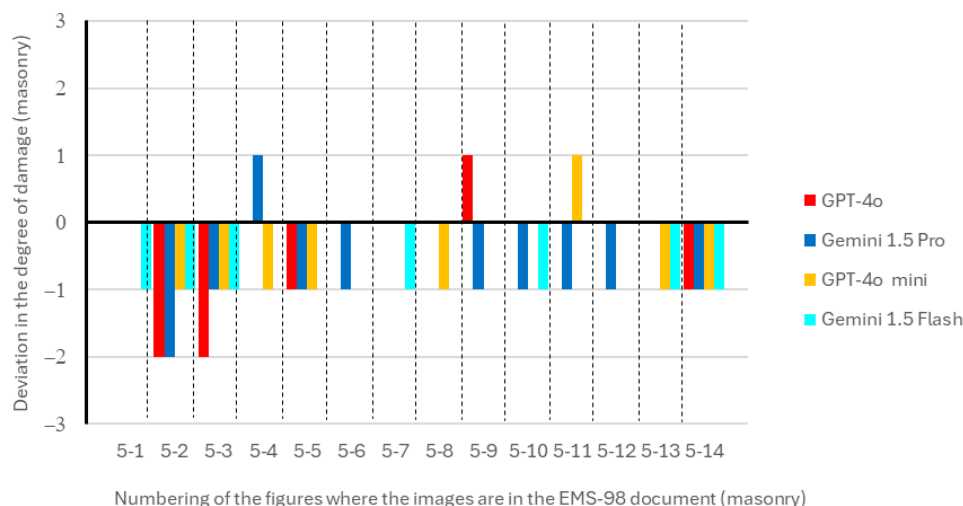


Figure 8. Deviations in the results obtained from the models tested on masonry buildings, compared to the damage levels presented in the EMS-98 document.

Table 2. Performance evaluation of GAI models for building earthquake damage assessment using EMS-98 imagery–RC buildings.

N.	Figure	EMS-98 Document	GPT-4o	Gemini 1.5 Pro	GPT-4o Mini	Gemini 1.5 Flash
1	5-16	D3	RC (90%) D3 (85%) ¹	RC (95%) D3 (85%)	RC (90%) D3 (85%)	RC (100%) D4 (100%)
2	5-17	D4	RC (90%) D4 (85%)	RC (95%) D5 (95%)	Masonry (90%) D3 (85%)	RC (100%) D4 (90%)
3	5-18	D5	RC (90%) D3 (85%)	RC (95%) D4 (85%)	RC (90%) D2 (85%)	RC (100%) D4 (90%)
4	5-19	D5	RC (90%) D5 (95%)	RC (95%) D5 (90%)	RC (85%) D5 (90%)	RC (100%) D4 (90%)
5	5-20	D4	RC (90%) D4 (85%)	RC (95%) D5 (90%)	RC (90%) D4 (85%)	RC (100%) D4 (90%)
6	5-21	D5	RC (90%) D5 (95%)	RC (95%) D5 (99%)	RC (85%) D5 (90%)	RC (100%) D4 (100%)
7	5-22	D5	RC (95%) D3 (90%)	RC (95%) D4 (85%)	RC (85%) D3 (80%)	RC (100%) D5 (100%)
8	5-23	D3	RC (90%) D3 (85%)	RC (95%) D2 (80%)	RC (85%) D2 (80%)	RC (100%) D3 (90%)

¹ In brackets is the degree of confidence that was presented by the model.

Accuracy rates for masonry buildings are as follows: 64.3% with GPT-4o; 28.6% with Gemini 1.5 Pro; 42.9% with GPT-4o mini; and 50.0% with Gemini 1.5 Flash. The mean errors in damage degree for masonry buildings are as follows: 0.50 with GPT-4o; 0.79 with Gemini 1.5 Pro; 0.57 with GPT-4o mini; and 0.50 with Gemini 1.5 Flash.

Accuracy rates for reinforced concrete buildings are as follows: 75.0% with GPT-4o; 37.5% with Gemini 1.5 Pro; 50.0% with GPT-4o mini; and 50.0% with Gemini 1.5 Flash. The mean errors in damage degree for reinforced concrete buildings are as follows: 0.50 with GPT-4o; 0.63 with Gemini 1.5 Pro; 0.88 with GPT-4o mini; and 0.50 with Gemini 1.5 Flash.

The use of RAG was also tested with the larger models (GPT-4o and Gemini 1.5 Pro), utilizing the official EMS-98 document to improve the results. In this context, both models were able to correctly identify the damage degree for all the images previously tested without RAG, due to the memorization effect of the images.

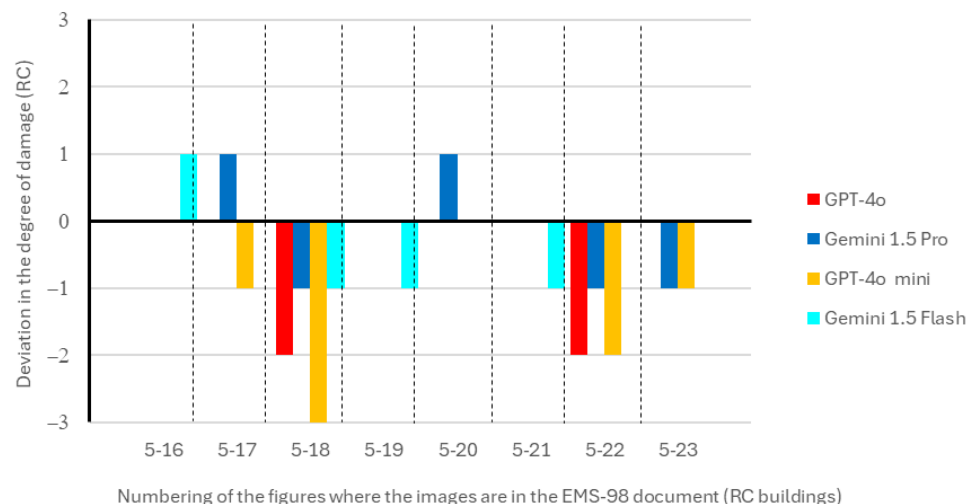


Figure 9. Deviations in the results obtained from the models tested on RC buildings, compared to the damage levels presented in the EMS-98 document.

4.5. Theoretical and Practical Implications

The theoretical implications of this study revolve around advancing the field of AI for earthquake damage assessment by integrating multimodal GAI models. These models show potential in replicating human-like decision-making processes, moving beyond traditional pattern recognition towards a more nuanced understanding of structural damage (a claim supported by the models' responses, as illustrated in Appendix C). This advancement lays a foundation for future developments in AI, particularly in enhancing model robustness and adaptability to diverse post-earthquake scenarios.

From a practical perspective, incorporating GAI models into damage assessment frameworks could drastically improve response times, especially in high-risk earthquake zones. By automating preliminary assessments, AI can reduce the burden on human evaluators and potentially operate in environments where human safety is at risk. However, for real-world deployment, larger datasets and thorough validation across diverse geographical and structural contexts are essential to ensure that these models can reliably function under varied conditions, considering the complexity of damage patterns in different regions.

5. Discussion

GAI models have been advancing at an accelerating rate. As a result, while some of the most recent GAI models were tested in this study, it is possible that they will already be outdated by the time this article is published. Despite this rapid advancement, the results obtained in this research can serve as a preliminary basis for the future development of software designed to perform automatic damage assessments of buildings. This could be especially important in earthquake-prone regions with a historical record of destructive earthquakes. This type of software could be crucial in regions like the Algarve (Portugal), which has been significantly impacted by earthquakes over time, most notably the one in 1755, where this region experienced the highest seismic intensities in Portugal [79].

It is important to emphasize that the results presented are those obtained from a single run of each of the tested models. Given the stochastic nature of these models, as described at the beginning of this article, it is possible that a slightly different outcome could be obtained in another run of the model. However, the impact of this randomness on the results is expected to be captured within the overall context of all evaluations, in principle.

It is noteworthy to observe the already significant capability of these models in identifying the existing structural systems in damaged buildings, whose images were analyzed by the GAI models. Only one error occurred across all tests conducted, as can be observed from the results listed in Tables 1 and 2.

Regarding the degree of damage, it is advisable to analyze the results obtained from masonry buildings separately from those with RC structures.

When examining the errors presented in Figure 8, for masonry buildings, it can be noted that these errors are generally within the range of one degree (+1 or −1), except for the buildings shown in Figure 5-2 and 5-3 of the EMS-98 document. The most significant errors (−2) are likely due to the fact that these images are in black and white in the EMS-98 document. Furthermore, most of the images are old and of low resolution, which may or may not present greater challenges for the models, particularly in the precise identification of small cracks. It is important to highlight that some tests were conducted with more recent images, whose results were not presented in this study for the reasons mentioned in Section 4.1, and the consistency of the results presented in this study was maintained. The variation of +1 or −1 observed in the damage degree indicates that the models still have some difficulty in accurately predicting the correct value for masonry buildings, but they are already very close. This type of variation might also be expected in the context of human assessment, which makes these results very promising from the outset.

Regarding buildings with RC structures, although the average error was lower than that obtained for masonry buildings, the maximum degree of error was higher (−3), as presented in Figure 9, which somewhat contradicted initial expectations, and there is no clear explanation for this fact.

In both types of buildings, the GPT-4o model demonstrated the best performance.

It also became evident that larger models do not always yield better results. For example, in this study, the Gemini 1.5 Flash model (more recent) delivered better results than the Gemini 1.5 Pro model (bigger but older).

To illustrate the type of responses generated by the models, several results were reproduced in Appendix C. These results include those obtained for Figure 5-4 (a correct identification of the damage level of a masonry building), Figure 5-17 (an incorrect identification of the building type), and Figure 5-19 (a correct identification of the damage level of a RC structure) from the EMS-98 document.

It is likely that, with a fine-tuning of the models or even with a simpler process such as the application of RAG techniques (as was tested, where the Gemini 1.5 Pro excelled), significantly higher levels of accuracy can be achieved in the future.

6. Conclusions

This study demonstrates the feasibility of utilizing Generative Artificial Intelligence (GAI) models for automated post-earthquake building damage assessment. Several key insights were derived from the experiments conducted:

- (1) **Feasibility of GAI Integration:** the integration of GAI models into software for post-earthquake damage assessment is viable and offers practical solutions for real-time structural evaluations.
- (2) **Variability in Model Accuracy:** when testing different GAI models, such as GPT-4o, GPT-4o mini, Gemini 1.5 Pro, and Gemini 1.5 Flash, the accuracy of damage assessment varied considerably. The overall accuracy rates were 68.2% for GPT-4o, 31.8% for Gemini 1.5 Pro, 45.5% for GPT-4o mini, and 50.0% for Gemini 1.5 Flash. Among these rates, GPT-4o exhibited the highest accuracy and the lowest average error margin.
- (3) **Implications for Future Use:** although the current performance of these models does not reach optimal levels, the application of advanced techniques like Retrieval-Augmented Generation (RAG) or model fine-tuning could substantially enhance their accuracy. With further improvements, GAI models hold significant potential for preliminary assessments of structural damage in buildings following seismic events.
- (4) **Need for Further Validation:** to ensure robustness at scale, further validation with larger datasets and real-world testing is essential, particularly in diverse earthquake scenarios.

Funding: This research was funded by Fundação para a Ciência e Tecnologia grant number LA/P/0069/2020 awarded to the ARNET (<https://doi.org/10.54499/LA/P/0069/2020>) and UID/00350/2020 awarded to CIMA-UAIG (<https://doi.org/10.54499/UIDB/00350/2020>).

Data Availability Statement: Data are contained within the article.

Acknowledgments: I wish to express my gratitude to OpenAI and Google for providing access to their AI models, which were instrumental in completing this study.

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

List of abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
CNN	Convolutional Neural Networks
DL	Deep Learning
EMS-98	European Macroseismic Scale 1998
FM	Foundation Model
GAI	Generative Artificial Intelligence
GAN	Generative Adversarial Networks
GPT	Generative Pretrained Transformer
LLM	Large Language Models
NLP	Natural Language Processing
NN	Neural Network
RAG	Retrieval-Augmented Generation
RC	Reinforced Concrete
RF	Random Forest
RL	Reinforcement Learning
SVM	Support Vector Machines
UAV	Unmanned Aerial Vehicles

Appendix A

The following is the Python code used to call the OpenAI API (all required libraries were previously installed on a personal computer) with the goal of extracting information from an image stored on the hard drive of a personal computer:

```
import os
import base64
import requests

# Function to read the contents of a file and return as a string
def read_file(path):
    with open(path, 'r', encoding='utf-8') as archive:
        return archive.read()

# Function to encode the image
def encode_image(image_path):
    with open(image_path, "rb") as image_file:
        return base64.b64encode(image_file.read()).decode('utf-8')

# Get the API key from the environment variable
api_key = os.getenv("OPENAI_API_KEY")

# Define the system headers
headers = {"Content-Type": "application/json", "Authorization": f"Bearer {api_key}"}

# Function to obtain results from analysing an image, using an OpenAI GPT model
```

```

def get_GPT_image(model, n, temperature, top_p, max_tokens, image_path, path_
prompt_image):

    # Get the base64 string of the image ['png', 'jpeg', 'gif', 'webp']
    base64_image = encode_image(image_path)

    # Reading the prompt
    image_prompt = read_file(path_prompt_image)

    # Defines the payload for the request
    payload = {
        "model": model,
        "messages": [
            {"role": "user", "content": [{"type": "text", "text": image_prompt},
            {"type": "image_url", "image_url":
                {"url": f"data:image/jpeg;base64,{base64_image}"}]}],
        "n": n, # Requests n answers at the prompt
        "temperature": temperature, # Adjust the temperature
        "top_p": top_p, # Adjust the top-P as required
        "max_tokens": max_tokens # Establishes the maximum number of tokens
    }

    # Makes the POST request to the OpenAI API
    response = requests.post(
"https://api.openai.com/v1/chat/completions", headers=headers, json=payload)

    # Checks if the request was successful
    if response.status_code == 200:
        completion = response.json()
    else:
        completion = response.text

    # Returns the completion object
    return completion

```

Appendix B

The following are the prompts adopted in the study using Markdown (which is a simple and easy-to-use markup language for formatting documents):

```

# Task
1. Look at the images of one building very carefully;
2. Identify the type of building based of the main structural system (masonry, rein-
forced concrete, timber or steel);
3. Classify the degree of seismic damage to the building;
4. Present your degree of confidence in the assessment you made (0 to 100%).
# Role
You are an earthquake engineer
# Format
## Material
1. Type of building
2. Degree of confidence
## Damage
1. Grade
2. Degree of confidence
## Explanation

```

Material
 ### Damage
 ### Confidence
 # Conditions

The degree of damage is assigned according to EMS-98 (D0, D1, D2, D3, D4 or D5).

Appendix C

The following is an example of a correct output of a model for Figure 5-4 of the EMS-98 document (GPT-4o):

Material
 Type of building: Masonry
 Degree of confidence: 90%
 Damage
 Grade: D4 (Very heavy damage)
 Degree of confidence: 85%

Explanation

Material

The building in the images appears to be made primarily of masonry. This assessment is based on the visible cracking patterns and the type of material exposed in the damaged areas, which resemble typical masonry construction with visible stones or bricks and mortar. The absence of visible reinforcing bars and the nature of the wall construction further support this conclusion.

Damage

The observed damage in the images includes:

Severe cracking and partial collapse of the walls.

Significant loss of structural integrity, particularly at the corners and edges of the building.

Large sections of the wall have crumbled and fallen away, exposing the interior.

Extensive damage to load-bearing walls, which indicates a compromise in the building's structural stability.

These observations align closely with the description of D4 (very heavy damage), which involves heavy structural damage and potentially very heavy non-structural damage. The building appears to be on the verge of collapse, which is characteristic of D4 damage, but it is not entirely demolished, distinguishing it from D5.

Confidence

The confidence levels are based on the clarity of the visible damage and the typical characteristics of masonry buildings. The visual evidence strongly supports the identification of the material and the classification of the damage, leading to high confidence levels in both assessments.

The following is an example of an incorrect output of a model for Figure 5-17 of the EMS-98 document (GPT-4o mini):

Material
 1. Type of building: Masonry
 2. Degree of confidence: 90%
 Damage
 1. Grade: D3
 2. Degree of confidence: 85%

Explanation

Material

The building appears to be constructed of masonry due to the visible brick and concrete elements that are typical in this type of construction. The presence of cracked walls and structural damage indicates that it is not reinforced concrete or steel.

Damage

The damage observed includes significant structural failure, including collapsed sections and large cracks, which aligns with a grade of D3 according to EMS-98. This grade indicates that the building has sustained considerable damage, affecting its integrity but not resulting in total collapse.

Confidence

The confidence levels are high due to clear visual indicators of material type and the extent of the damage observed, albeit with some uncertainty due to the limitations of assessing from a single image.

The following is an example of a correct output of a model for Figure 5-19 of the EMS-98 document (GPT-4o):

Material

Type of building: Reinforced concrete

Degree of confidence: 90%

Damage

Grade: D5

Degree of confidence: 95%

Explanation

Material

The building in the image appears to be constructed primarily of reinforced concrete. This assessment is based on the visible structure, which includes concrete floors and columns that are characteristic of reinforced concrete buildings. Additionally, the manner in which the building has collapsed suggests that the primary structural system is reinforced concrete, as masonry buildings typically collapse differently, and timber or steel structures would not exhibit the same kind of damage.

Damage

The damage observed in the building is extensive. The building has suffered partial collapse, with significant portions of the structure having disintegrated and fallen. This level of destruction corresponds to a Grade D5 classification according to the EMS-98 scale, which indicates very heavy damage or destruction. The visible signs include:

Complete collapse of large sections of the building.

Structural elements, such as columns and beams, have failed.

Extensive debris from the building's collapse is present.

Confidence

The degree of confidence in the material identification is 90% because the visible evidence strongly indicates reinforced concrete, though there might be some uncertainty due to the lack of close-up views of the construction materials. The confidence in the damage assessment is 95% because the extent of the collapse and the type of damage observed clearly match the criteria for Grade D5 damage as per EMS-98 standards.

References

1. PSBD. 2023 *Kahramanmaraş and Hatay Earthquakes Report*; Presidential Strategy and Budget Directorate: Ankara, Turkey, 2023.
2. Altunsu, E.; Güneş, O.; Öztürk, S.; Sorosh, S.; Sarı, A.; Beeson, S.T. Investigating the structural damage in Hatay province after Kahramanmaraş-Türkiye earthquake sequences. *Eng. Fail. Anal.* **2024**, *157*, 107857. [[CrossRef](#)]
3. Dasar, A.; Patah, D.; Caronge, M.A.; Apriansyah. Technical issues on forensic assessment of post-earthquake building damage—Case study of the Hajj Dormitory Building, Mamuju, Indonesia. *Structures* **2024**, *66*, 106772. [[CrossRef](#)]
4. Işık, E. Structural Failures of Adobe Buildings during the February 2023 Kahramanmaraş (Türkiye) Earthquakes. *Appl. Sci.* **2023**, *13*, 8937. [[CrossRef](#)]

5. Oforu-Ampong, K. Artificial intelligence research: A review on dominant themes, methods, frameworks and future research directions. *Telemat. Inform. Rep.* **2024**, *14*, 100127. [[CrossRef](#)]
6. Yenduri, G.; Ramalingam, M.; Selvi, G.C.; Supriya, Y.; Srivastava, G.; Maddikunta, P.K.R.; Raj, G.D.; Jhaveri, R.H.; Prabadevi, B.; Wang, W.; et al. GPT (Generative Pre-Trained Transformer)—A Comprehensive Review on Enabling Technologies, Potential Applications, Emerging Challenges, and Future Directions. *IEEE Access* **2024**, *12*, 54608–54649. [[CrossRef](#)]
7. Sarker, I.H. Machine Learning for Intelligent Data Analysis and Automation in Cybersecurity: Current and Future Prospects. *Ann. Data Sci.* **2023**, *10*, 1473–1498. [[CrossRef](#)]
8. Estêvão, J.M.C. Modelo Computacional de Avaliação do Risco Sísmico de Edifícios (Computational Model for Seismic Risk Assessment of Buildings). Master's Thesis, Instituto Superior Técnico, UTL, Lisbon, Portugal, 1998.
9. Klein, D. Neural Networks for Chess. The magic of deep and reinforcement learning revealed. *arXiv* **2022**, arXiv:2209.01506. [[CrossRef](#)]
10. Estêvão, J.M.C. Feasibility of using neural networks to obtain simplified capacity curves for seismic assessment. *Buildings* **2018**, *8*, 151. [[CrossRef](#)]
11. de-Miguel-Rodríguez, J.; Morales-Esteban, A.; Requena-García-Cruz, M.-V.; Zapico-Blanco, B.; Segovia-Verjel, M.-L.; Romero-Sánchez, E.; Estêvão, J.M.C. Fast Seismic Assessment of Built Urban Areas with the Accuracy of Mechanical Methods Using a Feedforward Neural Network. *Sustainability* **2022**, *14*, 5274. [[CrossRef](#)]
12. Estêvão, J.M.C.; Morales-Esteban, A.; Sá, L.F.; Ferreira, M.A.; Tomás, B.; Esteves, C.; Barreto, V.; Carreira, A.; Braga, A.; Requena-García-Cruz, M.-V.; et al. Improving the Earthquake Resilience of Primary Schools in the Border Regions of Neighbouring Countries. *Sustainability* **2022**, *14*, 15976. [[CrossRef](#)]
13. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet classification with deep convolutional neural networks. *Commun. ACM* **2017**, *60*, 84–90. [[CrossRef](#)]
14. LeCun, Y.; Boser, B.; Denker, J.S.; Henderson, D.; Howard, R.E.; Hubbard, W.; Jackel, L.D. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Comput.* **1989**, *1*, 541–551. [[CrossRef](#)]
15. Alom, M.Z.; Taha, T.M.; Yakopcic, C.; Westberg, S.; Sidike, P.; Nasrin, M.S.; Esesn, B.C.V.; Awwal, A.A.S.; Asari, V.K. The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches. *arXiv* **2018**, arXiv:1803.01164v2. [[CrossRef](#)]
16. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the Advances in Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA, 4–9 December 2017; pp. 5998–6008.
17. Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I. Improving Language Understanding by Generative Pre-Training. Available online: https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf (accessed on 10 July 2024).
18. Yildirim, I.; Paul, L.A. From task structures to world models: What do LLMs know? *Trends Cogn. Sci.* **2024**, *28*, 404–415. [[CrossRef](#)] [[PubMed](#)]
19. Matin, S.S.; Pradhan, B. Challenges and limitations of earthquake-induced building damage mapping techniques using remote sensing images—A systematic review. *Geocarto Int.* **2022**, *37*, 6186–6212. [[CrossRef](#)]
20. Ishiwatari, M. Leveraging Drones for Effective Disaster Management: A Comprehensive Analysis of the 2024 Noto Peninsula Earthquake Case in Japan. *Prog. Disaster Sci.* **2024**, *23*, 100348. [[CrossRef](#)]
21. Yu, R.; Li, P.; Shan, J.; Zhang, Y.; Dong, Y. Multi-feature driven rapid inspection of earthquake-induced damage on building facades using UAV-derived point cloud. *Measurement* **2024**, *232*, 114679. [[CrossRef](#)]
22. Yu, X.; Hu, X.; Song, Y.; Xu, S.; Li, X.; Song, X.; Fan, X.; Wang, F. Intelligent assessment of building damage of 2023 Turkey-Syria Earthquake by multiple remote sensing approaches. *npj Nat. Hazards* **2024**, *1*, 3. [[CrossRef](#)]
23. Jia, D.; Liu, Y.; Zhang, L. A rapid evaluation method of the seismic damage to buildings based on UAV images. *Geomatica* **2024**, *76*, 100006. [[CrossRef](#)]
24. Albahri, A.S.; Khaleel, Y.L.; Habeeb, M.A.; Ismael, R.D.; Hameed, Q.A.; Deveci, M.; Homod, R.Z.; Albahri, O.S.; Alamoodi, A.H.; Alzubaidi, L. A systematic review of trustworthy artificial intelligence applications in natural disasters. *Comput. Electr. Eng.* **2024**, *118*, 109409. [[CrossRef](#)]
25. Bhatta, S.; Dang, J. Seismic damage prediction of RC buildings using machine learning. *Earthq. Eng. Struct. Dyn.* **2023**, *52*, 3504–3527. [[CrossRef](#)]
26. Bhatta, S.; Dang, J. Machine Learning-Based Classification for Rapid Seismic Damage Assessment of Buildings at a Regional Scale. *J. Earthq. Eng.* **2024**, *28*, 1861–1891. [[CrossRef](#)]
27. Macchiarulo, V.; Giardina, G.; Miliello, P.; Aktas, Y.D.; Whitworth, M.R.Z. Integrating post-event very high resolution SAR imagery and machine learning for building-level earthquake damage assessment. *Bull. Earthq. Eng.* **2024**. [[CrossRef](#)]
28. Anniballe, R.; Noto, F.; Scalia, T.; Bignami, C.; Stramondo, S.; Chini, M.; Pierdicca, N. Earthquake damage mapping: An overall assessment of ground surveys and VHR image change detection after L'Aquila 2009 earthquake. *Remote Sens. Environ.* **2018**, *210*, 166–178. [[CrossRef](#)]
29. Adriano, B.; Yokoya, N.; Xia, J.; Miura, H.; Liu, W.; Matsuoka, M.; Koshimura, S. Learning from multimodal and multitemporal earth observation data for building damage mapping. *ISPRS J. Photogramm. Remote Sens.* **2021**, *175*, 132–143. [[CrossRef](#)]
30. Akhyar, A.; Asyraf Zulkifley, M.; Lee, J.; Song, T.; Han, J.; Cho, C.; Hyun, S.; Son, Y.; Hong, B.-W. Deep artificial intelligence applications for natural disaster management systems: A methodological review. *Ecol. Indic.* **2024**, *163*, 112067. [[CrossRef](#)]

31. Amanollah, H.; Asghari, A.; Mashayekhi, M.; Zahrai, S.M. Damage detection of structures based on wavelet analysis using improved AlexNet. *Structures* **2023**, *56*, 105019. [[CrossRef](#)]
32. Bai, Z.; Liu, T.; Zou, D.; Zhang, M.; Hu, Q.; Zhou, A.; Li, Y. Multi-scale image-based damage recognition and assessment for reinforced concrete structures in post-earthquake emergency response. *Eng. Struct.* **2024**, *314*, 118402. [[CrossRef](#)]
33. Bhatta, S.; Dang, J. Multiclass seismic damage detection of buildings using quantum convolutional neural network. *Comput. -Aided Civ. Infrastruct. Eng.* **2024**, *39*, 406–423. [[CrossRef](#)]
34. Braik, A.M.; Koliou, M. Automated building damage assessment and large-scale mapping by integrating satellite imagery, GIS, and deep learning. *Comput. -Aided Civ. Infrastruct. Eng.* **2024**, *39*, 2389–2404. [[CrossRef](#)]
35. Cheng, M.-Y.; Khasani, R.R.; Citra, R.J. Image-based preliminary emergency assessment of damaged buildings after earthquake: Taiwan case studies. *Eng. Appl. Artif. Intell.* **2023**, *126*, 107164. [[CrossRef](#)]
36. Kijewski-Correa, T.; Canales, E.; Hamburger, R.; Lochhead, M.; Mbabazi, A.; Presuma, L. A hybrid model for post-earthquake performance assessments in challenging contexts. *Bull. Earthq. Eng.* **2024**. [[CrossRef](#)]
37. Matin, S.S.; Pradhan, B. Earthquake-Induced Building-Damage Mapping Using Explainable AI (XAI). *Sensors* **2021**, *21*, 4489. [[CrossRef](#)] [[PubMed](#)]
38. Ogunjinmi, P.D.; Park, S.-S.; Kim, B.; Lee, D.-E. Rapid Post-Earthquake Structural Damage Assessment Using Convolutional Neural Networks and Transfer Learning. *Sensors* **2022**, *22*, 3471. [[CrossRef](#)] [[PubMed](#)]
39. Sublime, J.; Kalinicheva, E. Automatic Post-Disaster Damage Mapping Using Deep-Learning Techniques for Change Detection: Case Study of the Tohoku Tsunami. *Remote Sens.* **2019**, *11*, 1123. [[CrossRef](#)]
40. Xia, H.; Wu, J.; Yao, J.; Zhu, H.; Gong, A.; Yang, J.; Hu, L.; Mo, F. A Deep Learning Application for Building Damage Assessment Using Ultra-High-Resolution Remote Sensing Imagery in Turkey Earthquake. *Int. J. Disaster Risk Sci.* **2023**, *14*, 947–962. [[CrossRef](#)]
41. Yilmaz, M.; Dogan, G.; Arslan, M.H.; Ilki, A. Categorization of Post-Earthquake Damages in RC Structural Elements with Deep Learning Approach. *J. Earthq. Eng.* **2024**, *28*, 2620–2651. [[CrossRef](#)]
42. You, C.; Liu, W.; Hou, L. Convolutional Neural Networks for Structural Damage Identification in Assembled Buildings. *Math. Probl. Eng.* **2022**, *2022*, 2326903. [[CrossRef](#)]
43. Zhan, Y.; Liu, W.; Maruyama, Y. Damaged Building Extraction Using Modified Mask R-CNN Model Using Post-Event Aerial Images of the 2016 Kumamoto Earthquake. *Remote Sens.* **2022**, *14*, 1002. [[CrossRef](#)]
44. Marano, G.C.; Rosso, M.M.; Aloisio, A.; Cirrincione, G. Generative adversarial networks review in earthquake-related engineering fields. *Bull. Earthq. Eng.* **2024**, *22*, 3511–3562. [[CrossRef](#)]
45. Cheng, M.-Y.; Sholeh, M.N.; Kwek, A. Computer vision-based post-earthquake inspections for building safety assessment. *J. Build. Eng.* **2024**, *94*, 109909. [[CrossRef](#)]
46. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; Arx, S.v.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* **2022**, arXiv:2108.07258v3. [[CrossRef](#)]
47. LlamaTeam. The Llama 3 Herd of Models. Available online: <https://ai.meta.com/research/publications/the-llama-3-herd-of-models/> (accessed on 24 July 2024).
48. Eloundou, T.; Manning, S.; Mishkin, P.; Rock, D. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv* **2023**, arXiv:2303.10130v5. [[CrossRef](#)]
49. Betker, J.; Goh, G.; Jing, L.; Brooks, T.; Wang, J.; Li, L.; Ouyang, L.; Zhuang, J.; Lee, J.; Guo, Y.; et al. Improving Image Generation with Better Captions. Available online: <https://cdn.openai.com/papers/dall-e-3.pdf> (accessed on 12 July 2024).
50. OpenAI. GPT-4 Technical Report. *arXiv* **2023**, arXiv:2303.08774v6. [[CrossRef](#)]
51. Ray, P.P. ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet Things Cyber Phys. Syst.* **2023**, *3*, 121–154. [[CrossRef](#)]
52. Ooi, K.-B.; Tan, G.W.-H.; Al-Emran, M.; Al-Sharafi, M.A.; Capatina, A.; Chakraborty, A.; Dwivedi, Y.K.; Huang, T.-L.; Kar, A.K.; Lee, V.-H.; et al. The Potential of Generative Artificial Intelligence Across Disciplines: Perspectives and Future Directions. *J. Comput. Inf. Syst.* **2023**, 1–32. [[CrossRef](#)]
53. OpenAI. Dall-E 3. Available online: <https://openai.com/index/dall-e-3/> (accessed on 12 September 2024).
54. Yao, Y.; Duan, J.; Xu, K.; Cai, Y.; Sun, Z.; Zhang, Y. A survey on large language model (LLM) security and privacy: The Good, The Bad, and The Ugly. *High Confid. Comput.* **2024**, *4*, 100211. [[CrossRef](#)]
55. OpenAI. Tokenizer. Available online: <https://platform.openai.com/tokenizer> (accessed on 28 July 2024).
56. ChameleonTeam. Chameleon: Mixed-Modal Early-Fusion Foundation Models. *arXiv* **2024**, arXiv:2405.09818v1. [[CrossRef](#)]
57. Neelakantan, A.; Xu, T.; Puri, R.; Radford, A.; Han, J.M.; Tworek, J.; Yuan, Q.; Tezak, N.; Kim, J.W.; Hallacy, C.; et al. Text and Code Embeddings by Contrastive Pre-Training. *arXiv* **2022**, arXiv:2201.10005v1. [[CrossRef](#)]
58. Sanderson, G. But what is a GPT? Visual intro to transformers | Chapter 5, Deep Learning. Available online: <https://www.3blue1brown.com/lessons/gpt> (accessed on 17 July 2024).
59. Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.; Dean, J. Distributed Representations of Words and Phrases and their Compositionality. *arXiv* **2013**, arXiv:1310.4546v1. [[CrossRef](#)]
60. Li, Y.; Xu, L.; Tian, F.; Jiang, L.; Zhong, X.; Chen, E. Word Embedding Revisited: A New Representation Learning and Explicit Matrix Factorization Perspective. In Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, Buenos Aires, Argentina, 25–31 July 2015; pp. 3650–3656.

61. Sanderson, G. Visualizing Attention, a Transformer’s Heart | Chapter 6, Deep Learning. Available online: <https://www.3blue1brown.com/lessons/attention> (accessed on 17 July 2024).
62. Gemini Team, G. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv* **2024**, arXiv:2403.05530v3. [[CrossRef](#)]
63. Holtzman, A.; Buys, J.; Du, L.; Forbes, M.; Choi, Y. The Curious Case of Neural Text Degeneration. *arXiv* **2020**, arXiv:1904.09751v2. [[CrossRef](#)]
64. White, J.; Fu, Q.; Hays, S.; Sandborn, M.; Olea, C.; Gilbert, H.; Elnashar, A.; Spencer-Smith, J.; Schmidt, D.C. A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT. *arXiv* **2023**, arXiv:2302.11382v1. [[CrossRef](#)]
65. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* **2023**, *55*, 248. [[CrossRef](#)]
66. López Espejel, J.; Ettifouri, E.H.; Yahaya Alassan, M.S.; Chouham, E.M.; Dahhane, W. GPT-3.5, GPT-4, or BARD? Evaluating LLMs reasoning ability in zero-shot setting and performance boosting through prompts. *Nat. Lang. Process. J.* **2023**, *5*, 100032. [[CrossRef](#)]
67. Lo, L.S. The Art and Science of Prompt Engineering: A New Literacy in the Information Age. *Internet Ref. Serv. Q.* **2023**, *27*, 1–8. [[CrossRef](#)]
68. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. In Proceedings of the 34th International Conference on Neural Information Processing Systems, Vancouver, BC, Canada, 6–12 December 2020; p. 159.
69. Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.L.; Cao, Y.; Narasimhan, K. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *arXiv* **2023**, arXiv:2305.10601v2. [[CrossRef](#)]
70. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv* **2019**, arXiv:1810.04805v2. [[CrossRef](#)]
71. Wu, S.; Xiong, Y.; Cui, Y.; Wu, H.; Chen, C.; Yuan, Y.; Huang, L.; Liu, X.; Kuo, T.-W.; Guan, N.; et al. Retrieval-Augmented Generation for Natural Language Processing: A Survey. *arXiv* **2023**, arXiv:2407.13193v2. [[CrossRef](#)]
72. Pu, Y.; He, Z.; Qiu, T.; Wu, H.; Yu, B. Customized Retrieval Augmented Generation and Benchmarking for EDA Tool Documentation QA. *arXiv* **2024**, arXiv:2407.15353v1.
73. Arasteh, S.T.; Lotfinia, M.; Bressemer, K.; Siepmann, R.; Ferber, D.; Kuhl, C.; Kather, J.N.; Nebelung, S.; Truhn, D. RadioRAG: Factual Large Language Models for Enhanced Diagnostics in Radiology Using Dynamic Retrieval Augmented Generation. *arXiv* **2024**, arXiv:2407.15621v1. [[CrossRef](#)]
74. Nicodemo, G.; Pittore, M.; Masi, A.; Manfredi, V. Modelling exposure and vulnerability from post-earthquake survey data with risk-oriented taxonomies: AeDES form, GEM taxonomy and EMS-98 typologies. *Int. J. Disaster Risk Reduct.* **2020**, *50*, 101894. [[CrossRef](#)]
75. Grünthal, G. *European Macroseismic Scale 1998. Volume 15.*; Centre Européen de Géodynamique et de Séismologie: Luxembourg, 1998; p. 99.
76. OpenAI. GPT-4o. Available online: <https://platform.openai.com/playground/chat?models=gpt-4o> (accessed on 30 July 2024).
77. Google. Google AI Studio. Available online: https://aistudio.google.com/app/prompts/new_chat (accessed on 30 July 2024).
78. OpenAI. GPT-4o mini. Available online: <https://platform.openai.com/playground/chat?models=gpt-4o-mini> (accessed on 30 July 2024).
79. Chester, D.K.; Chester, O.K. The impact of eighteenth century earthquakes on the Algarve region, southern Portugal. *Geogr. J.* **2010**, *176*, 350–370. [[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.