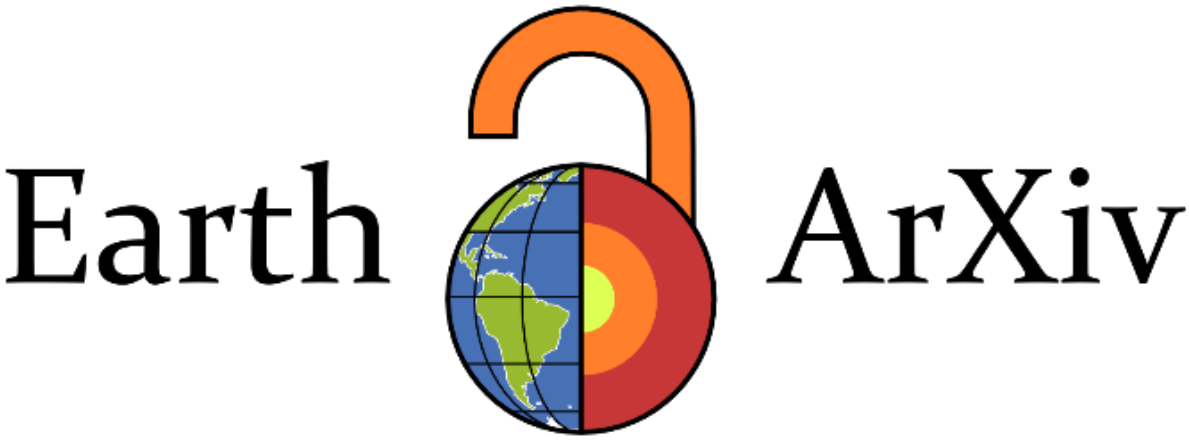

Operationalising EMS-98 Damage Classigication: A UAV-to-GIS Pipeline for Macroseismic Survey Support

- *Giovanni Galli - Giovanni.galli@tuta.com. University of Bologna;*
- *Marco Dubbini - Marco.dubbini@unibo.it. University of Bologna;*
- *Filippo Bernardini - Filippo.bernardini@ingv.it . Istituto Nazionale di Geofisica e Vulcanologia*
- *Luca Arcoraci - Luca.arcoraci@ingv.it. Istituto Nazionale di Geofisica e Vulcanologia*

Article submitted for review in the Elsevier International Journal of Disaster Risk Reduction

1. Introduction

Rapid and reliable assessment of damage to the built environment is a critical component of disaster risk reduction (DRR), particularly in the aftermath of earthquakes, where timely information on the extent and severity of structural damage underpins decisions related to safety evaluation, resource allocation, and recovery planning (Matin and Pradhan, 2022; Al Shafian and Hu, 2024). During the recovery phase of a seismic emergency, damage assessment activities provide the basis for prioritising inspections, coordinating technical interventions, and quantifying macroseismic intensity at the locality level. Within European practice, this latter task is formalised through macroseismic surveying, in which building damage is systematically observed and graded according to the European Macroseismic Scale 1998 (EMS-98) to derive an intensity value for each affected locality (Grünthal, 1998; Tertulliani et al., 2025).

Despite its operational importance, post-earthquake damage assessment remains predominantly reliant on ground-based visual inspections conducted by trained personnel (Tertulliani et al., 2011; Jozi et al., 2024). While robust and methodologically well established, such approaches are inherently constrained by limited scalability, high operational costs, and reduced accessibility in severely affected areas, often resulting in delays that can hinder effective recovery and impede the spatial completeness of macroseismic datasets (Tertulliani et al., 2011; Jozi et al., 2024). The documented disparity between the number of affected structures and the time required for systematic building-by-building inspection is a recurring limitation reported across recent seismic events, contributing to survey coverage that is concentrated in severely affected areas and sparse elsewhere *(Rossi et al., 2019)*.

Remote sensing technologies have been increasingly adopted to address these limitations. Satellite imagery enables rapid, large-scale assessments but remains constrained by a nadir viewing geometry, limited spatial resolution for small damage manifestations, and susceptibility to atmospheric conditions, all of which restrict its ability to capture façade-level damage and fine-grained structural details (Brunner et al., 2010; Ainscoe et al., 2025). In contrast, Unmanned Aerial Vehicles (UAVs) provide high-resolution, multi-view imagery, enabling detailed observation of building façades and structural elements. Their flexibility, low cost, and capacity to operate in hazardous or inaccessible environments make them particularly suitable for post-disaster data acquisition at the scale and detail required for automated damage detection and

classification (Fernandez Galarreta et al., 2015; Kerle et al., 2019; Khan et al., 2022; Tavakol Sadrabadi et al., 2024).

In parallel, computer vision techniques have demonstrated strong performance in damage detection, classification, and segmentation tasks, and modern single-stage architectures have reached levels of accuracy and efficiency compatible with operational deployment (Bai et al., 2021; Zou et al., 2024). However, the majority of existing studies conceptualise these tasks as isolated problems, focusing on model performance on reference datasets rather than on how such model outputs can be operationalised within emergency management workflows (Al Shafian and Hu, 2024; Kallas and Napolitano, 2025). Consequently, a persistent gap remains between image-based damage detection and its translation into actionable information employable by the emergency actors involved in damage assessment activities.

This limitation is particularly evident in macroseismic assessment contexts. The EMS-98 relies on observed building damage to classify earthquake intensity and to inform recovery processes, yet existing AI-based approaches rarely attempt to map detected damage features to macroseismic damage grades, nor do they provide mechanisms to aggregate multiple damage observations into coherent building-level assessments (Gerke and Kerle, 2011; Calantropio et al., 2021). Moreover, most methods do not produce outputs in formats readily interoperable within Geographic Information Systems (GIS), limiting their integration into operational decision-making processes (Jozi et al., 2024; Kallas and Napolitano, 2025).

To address these gaps, this study presents the Macroseismic Survey Mapper (MSM), a GeoAI workflow designed to support post-earthquake macroseismic surveying. The MSM is conceptualised as an end-to-end pipeline that links four operational stages: (i) the seismic event, (ii) the UAV imagery acquisition mission, (iii) instance-level damage detection followed by building-level aggregation, and (iv) export of a geodataframe (GeoPackage). At the core of this pipeline sits the MSM model, the instance-segmentation component that detects façade damage manifestations and classifies each instance according to the five EMS-98 severity grades. Throughout the present work, the pipeline (the MSM) and its underlying detection model (the MSM model) are treated as conceptually distinct: the pipeline defines the operational workflow into which the model is embedded, while the model is the AI mechanism underpinning damage interpretation within that workflow.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature on remote sensing for damage assessment, deep learning for damage detection, and the operational integration of such methods with macroseismic frameworks. Section 3 describes the dataset, model architecture, training strategy, and pipeline design. Section 4 presents and discusses the experimental results. Section 5 concludes with implications for operational deployment and directions for future work.

2. Background

2.1 Remote sensing and UAV imagery for post-disaster damage assessment

Remote sensing has been extensively employed for post-disaster damage assessment, primarily leveraging satellite imagery for rapid, large-scale mapping. Approaches based on visual interpretation, change

detection, and, more recently, deep learning have demonstrated substantial effectiveness in identifying affected areas and providing initial situational awareness (Matin and Pradhan, 2022; Al Shafian and Hu, 2024). Very High Resolution (VHR) optical and Synthetic Aperture Radar (SAR) imagery, in particular, are well established for rapid mapping of severe damage and building collapse at regional scales (Brunner et al., 2010; Ainscoe et al., 2025). However, satellite-based methods remain fundamentally limited in their capacity to capture detailed structural damage. The vertical viewing geometry restricts observations to rooftops, while the spatial resolution constrains the detection of fine-scale damage features on building façades (Duarte et al., 2017).

UAV-based imaging has emerged as a complementary approach capable of overcoming these limitations. By enabling high-resolution, multi-angle data acquisition at centimetre-level ground sampling distances, UAVs facilitate detailed inspection of building façades and structural components. This capability is particularly relevant in dense urban environments and in scenarios where ground access is restricted by debris, structural instability, or safety perimeters (Khan et al., 2022; Tavakol Sadrabadi et al., 2024).

A growing body of work has demonstrated the value of directly processing oblique UAV imagery for façade damage assessment. Fernandez Galarreta et al. (2015) combined multi-view UAV imagery with object-based image analysis, showing that oblique UAV coverage yields damage evidence not interpretable from nadir perspectives, while Calantropio et al. (2021) applied UAV orthophotos to perform building-level damage classification at regional scale over the 2016 Central Italy earthquake area. More recent contributions have extended these directions. Jozi et al. (2024) developed a single-image, feature-based framework that combines Gray-Level Co-occurrence Matrix texture descriptors (dissimilarity and homogeneity) with novel edge-based indices in a Naïve Bayesian classifier, reaching 89.3% accuracy in discriminating damaged from undamaged residential buildings from UAV imagery and reaffirming the diagnostic value of side-view acquisitions. Cheng et al. (2024) coupled a region-based convolutional neural network with UAV photogrammetry to detect impacted infrastructure and quantify post-hurricane debris volumes in Southeast Texas, reporting an 81% average precision for debris detection and identifying an image overlap ratio of 60–70% as the optimal trade-off between geometric accuracy and acquisition efficiency for time-critical response.

While alternative pipelines that rely on LiDAR point clouds can deliver spatially precise, volumetric damage characterisations (Cheng et al., 2024; Kallas and Napolitano, 2025; Kim and Leite, 2026), they typically demand specialised sensors, strict acquisition protocols, and substantial storage and computational resources. The present study therefore adopts a data-light stance, operating directly on 2D façade imagery and foregoing any intermediate 3D reconstruction step, in line with the operational priorities of rapid post-earthquake damage assessment.

2.2 Deep learning for damage detection

The application of deep learning has significantly advanced automated damage detection in remote sensing imagery. Convolutional Neural Networks (CNNs) enable hierarchical feature extraction, supporting a range of computer vision tasks including image classification, object detection, and semantic or instance segmentation. Architectures such as YOLO that include Feature Pyramid Networks have been widely

adopted due to their ability to balance detection accuracy and computational efficiency, making them suitable candidates for operational deployment (Zou et al., 2024; Al Shafian and Hu, 2024). The single-stage design of YOLO architectures, in which detection is formulated as a single regression pass over the image, is particularly suited to time-critical UAV applications (Xing et al., 2022) and has been demonstrated to perform well on post-event imagery alone, without recourse to pre-event reference data (Zou et al., 2024). For damage detection tasks specifically, instance segmentation approaches, in which each damage instance is delineated by a pixel-level mask in addition to a class label, have proven superior to semantic segmentation in contexts where individual damage occurrences must be separately attributed, quantified, and exported (Bai et al., 2021; Kim and Cho, 2020).

Despite these advances, existing approaches exhibit several limitations. First, many studies focus on binary or coarse damage classification, providing limited insight into the nature and extent of structural damage (Bouchard et al., 2022; Wiguna et al., 2024). Second, a substantial portion of the literature addresses specific damage types in isolation without accounting for the coexistence and interaction of multiple damage manifestations within a single structure. For instance, Kim and Cho (2020) successfully train a Mask R-CNN to simultaneously segment cracks, efflorescence, rebar exposure, and spalling on reinforced-concrete elements, but do not aggregate these to a building-level structural-grading framework. As a result, current models often fail to capture the complexity of real-world damage patterns, limiting their applicability for comprehensive structural assessment within established grading frameworks. A related methodological challenge is the limited size and event-specific nature of post-earthquake training datasets. Transfer learning, in which model parameters are initialised from weights trained on a larger related dataset, has become a standard response to this constraint (Zhuang et al., 2020; Yang et al., 2021; Bouchard et al., 2022).

2.3 From damage detection to decision-ready, EMS-98-aligned data products

A persistent limitation across the damage detection literature is the disconnect between model outputs and their operational use. Emergency management workflows generally require geospatially explicit data products — vector layers or geodatabases that can be ingested into GIS environments for mapping, prioritisation, and the coordination of inspections — yet the majority of deep-learning damage-detection studies stop at image-level outputs and do not translate detections into such formats (Jozi et al., 2024; Wiguna et al., 2024; Kallas and Napolitano, 2025). Within the European context, several studies have attempted alignment with the EMS-98, the most widely adopted reference for post-earthquake damage grading and intensity assessment (Grünthal, 1998; Gerke and Kerle, 2011; Tertulliani et al., 2011; Rossi et al., 2019); however, translating its five-grade scheme into reliable automated classification has proven difficult. Extreme grades (negligible damage and total collapse) are visually distinctive and consistently classified, whereas intermediate grades exhibit overlapping visual cues that are hard to discriminate (Kakooei and Baleghi, 2017; Song et al., 2020; Calantropio et al., 2021). Indeed, distinguishing adjacent damage degrees in the intermediate range (grades 1, 2 and 3) is a challenge even for professional ground surveyors (Tertulliani et al., 2025). This difficulty has driven many authors toward reduced or binary class schemes, fragmenting the literature and undermining cross-study comparability.

A more fundamental shortcoming is that, even where damage is reliably detected, models rarely translate their outputs into structured, decision-ready information products. Kallas and Napolitano (2025) make this gap explicit: in a comparative review of twenty prior damage-detection studies, they show that none of the surveyed approaches — whether based on satellite and aerial imagery, close-up imagery, or 3D point clouds — produce an automated, structured data output suitable for downstream operational use. Most studies are limited to 2D classification of a single damage type — typically cracks on concrete or pavement — and those that incorporate 3D spatial information generally produce visualisations rather than data structures designed for downstream operational use. Their own approach addresses this gap by projecting Mask R-CNN segmentation masks (trained on 6,000 post-disaster images of historic masonry from the 2020 Beirut explosion) onto photogrammetric 3D point clouds and exporting clustered damage entities as a structured dataframe with class and geometric attributes. Jozi et al. (2024) raise a parallel concern, identifying the absence of end-to-end automation — from UAV image acquisition through to georeferenced damage records — as a principal barrier to the rapid operational deployment of AI-based assessment tools.

These limitations are particularly consequential for macroseismic intensity assessment. The EMS-98 has become the standard reference for post-earthquake damage assessment in European civil-protection workflows (Rossi et al., 2019; Tertulliani et al., 2011, 2025). Notably, the scale's original specification explicitly anticipated a basis for computerised macroseismic intensity evaluation (Grünthal, 1998), motivating the production of georeferenced data products consistent with EMS-98 building-level grading. In current practice, however, AI-based damage detections are rarely aggregated into building-level EMS-98 grades, nor delivered in formats compatible with civil-protection and macroseismic-survey workflows. The MSM pipeline developed in this study responds to these gaps and adopts an approach analogous to Kallas and Napolitano (2025) but tailored to macroseismic survey needs, delivering an EMS-98-graded geodataframe of the affected buildings, georeferenced through UAV GNSS and camera-orientation metadata and directly interoperable with GIS environments.

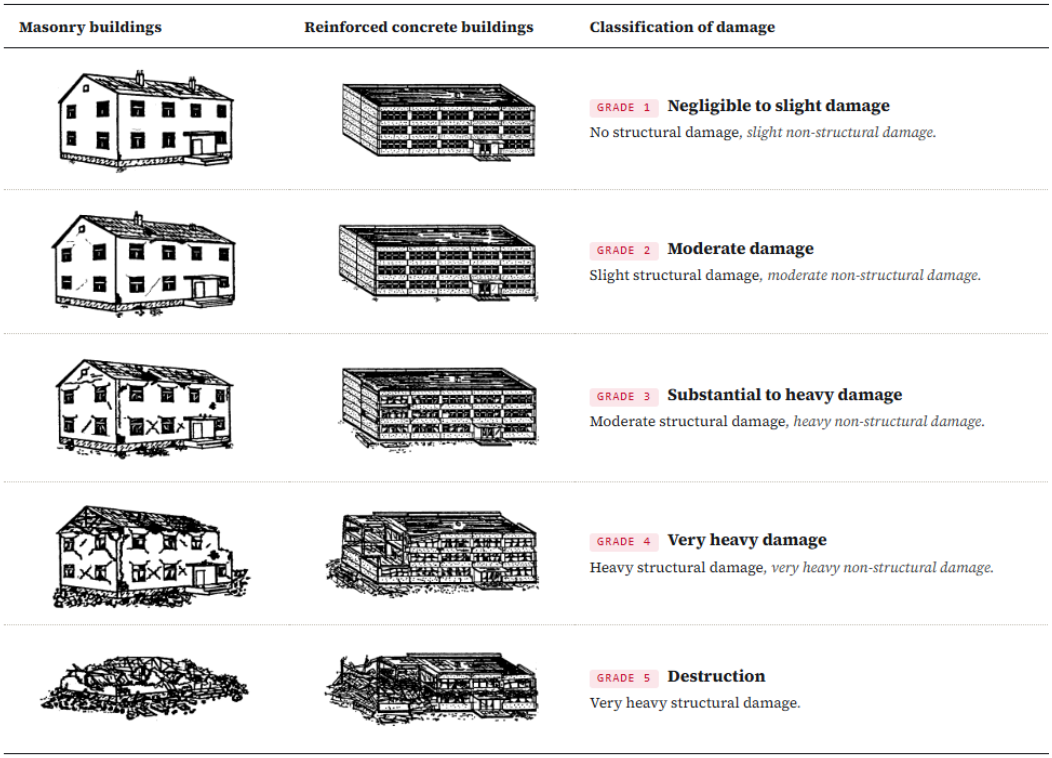

Figure 1. EMS-98 damage classification scheme for masonry and reinforced-concrete buildings (adapted from Grünthal, 1998).

2.4 Research gaps and contribution

Three gaps motivate the present work: (i) no existing pipeline links UAV imagery acquisition, automated EMS-98 damage classification, and export of a georeferenced building-level geodataframe within an end-to-end workflow; (ii) deep-learning models for EMS-98-aligned classification are typically trained on controlled or single-event datasets, limiting their generalisation; and (iii) model outputs are rarely structured for institutional use, providing neither GIS-ready formats nor per-building damage grades.

In response to these gaps, this study develops the Macroseismic Survey Mapper (MSM), conceived at two complementary levels. At the workflow level, the MSM is an end-to-end pipeline that links UAV imagery acquisition, automated damage detection and export of a GIS-interoperable geodataframe. At the model level, the MSM is built on an instance-segmentation network — the MSM model — that detects façade damage manifestations, classifies each instance across the five EMS-98 severity grades, and aggregates instance-level predictions into a single damage grade per building by assigning the highest detected grade.

As illustrated in Figure 2, the MSM pipeline is intended to be deployed on UAV-derived imagery in a four-stage sequence. After the onset of a seismic event, a UAV mission is planned over the affected area to collect an image set of building façades together with the EXIF metadata of each acquisition; the imagery is then passed through the MSM model, which detects damage manifestations and aggregates them into a single EMS-98 damage grade per building; finally, detections and associated EXIF metadata are projected

onto a horizontal plane and intersected, where available, with building footprints to produce a geodataframe of building locations and aggregated damage grades, exported as a GeoPackage layer.

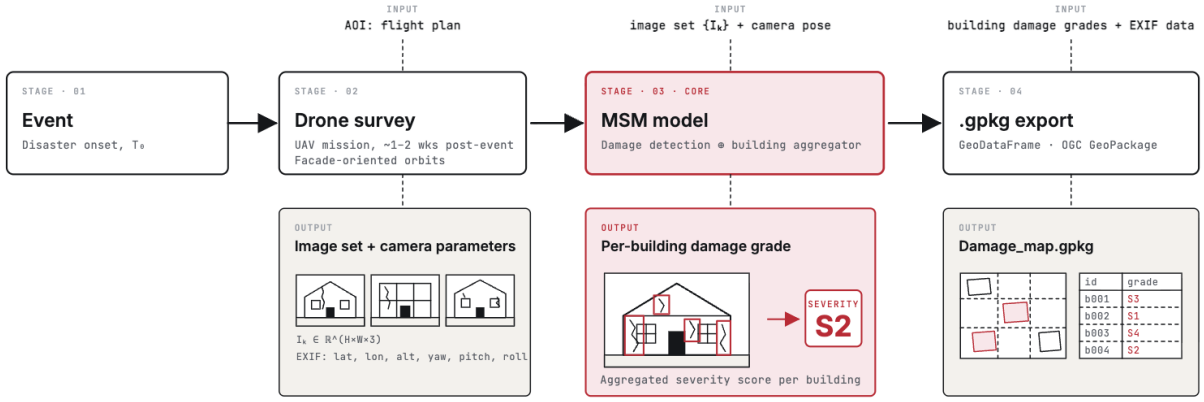

Figure 2. Conceptual overview of the end-to-end MSM pipeline. Stage 1 (event) and Stage 2 (UAV survey) define the operational inputs to the pipeline; Stage 3 is the core MSM model, which performs damage detection and building-level aggregation; Stage 4 exports the result as a GeoPackage (.gpkg) layer suitable for ingestion in GIS environments.

3. Materials and Methods

This study proposes a five-step methodology to develop and operationalise the MSM. The methodology comprises: (1) compilation of a heterogeneous training corpus from the INGV Database Fotografico Macrosismico (DFM); (2) instance-level annotation of damage manifestations across the EMS-98 severity scheme; (3) a transfer-learning strategy in which a YOLOv11 segmentation network is first pretrained on the Bai et al. (2021) crack-and-spalling dataset and then fine-tuned on the DFM imagery; (4) a training schedule with validation and test-set evaluation; and (5) operationalisation, in which instance-level predictions are aggregated into building-level EMS-98 grades and exported as a georeferenced geodataframe through UAV GNSS and camera-orientation metadata. The remainder of this section details each step.

3.1 Dataset

The training corpus was sourced from the DFM, a publicly accessible archive of ground-survey photographs collected by expert teams during past Italian seismic events. *Two Authors of the present study contributed directly to the original DFM acquisition campaigns, providing first-hand familiarity with the field acquisition conditions reflected in the corpus.* The DFM is opportunistic in nature: photographs were predominantly captured by human surveyors using DSLR cameras, which introduces strong variability in viewpoint and distance to the structure, illumination and shadowing, background clutter, and partial occlusion. Although challenging from a learning standpoint, this heterogeneity faithfully reflects the operational acquisition conditions of macroseismic surveys and was therefore retained as a defining property of the dataset.

Imagery was retrieved through a REST query service provided by the DFM. Six earthquake events spanning seismic activity from 2009 to 2016 and a moment-magnitude (Mw) range of 4.23–6.29 were targeted, yielding a corpus of 866 images (Table 1). From a deep-learning perspective this constitutes a low-data regime, in which training a high-capacity instance-segmentation network from scratch would be prone to overfitting and unstable convergence; the limitation is intrinsic to the opportunistic nature of the dataset and motivates the transfer-learning strategy described in Section 3.2.

Epicentral area	**Event date**	**M_w**	**No. of images**
Aquilano	2009-04-06	6.29	509
Valle del Tevere	2009-12-15	4.23	69
Pianura emiliana	2012-05-20	6.09	113
Pianura emiliana	2012-05-29	5.90	134
Lunigiana	2013-06-21	5.36	34
Valnerina	2016-10-26	6.07	7
Total			**866**

Table 1. Images retrieved from the INGV DFM by epicentral area, event date, and moment magnitude (Mw).

Although the EMS-98 distinguishes between masonry (M) and reinforced-concrete (RC) structural typologies, an initial fully disaggregated ten-class formulation produced fragmented sample sizes that compromised the learning of intermediate grades. The two construction-type taxonomies were therefore consolidated into a unified five-class severity scheme (S1–S5), a simplification consistent with strategies adopted in comparable studies (Kakooei and Baleghi, 2017; Song et al., 2020). An ablation study formally comparing the disaggregated 10-class formulation with the unified 5-class scheme — and quantifying the resulting effect on per-grade learnability — is reported in Section 4.1.1.

The corpus was partitioned into training, validation, and test subsets following a 70/20/10 split. A critical constraint was that the DFM frequently contains multiple photographs of the same building; to prevent leakage across subsets from near-duplicate views, all images of a single structure were assigned to the same partition. After splitting, the dataset comprised 587 training, 170 validation, and 109 test images, distributed across severity grades as reported in Table 2.

Damage level	**Train**	**Val**	**Test**	**Total**
S1 (M1+RC1)	27	8	9	44
S2 (M2+RC2)	166	48	28	242
S3 (M3+RC3)	200	56	33	289
S4 (M4+RC4)	119	37	24	180
S5 (M5+RC5)	75	21	15	111
Total	**587**	**170**	**109**	**866**

Table 2. Train/validation/test split for the MSM model, after merging masonry (M) and reinforced-concrete (RC) classes into unified severity levels (S1–S5).

To support transfer learning, training was initialised from a checkpoint derived from the dataset curated by Bai et al. (2021), comprising 2,021 labelled images of cracks and spalling across three spatial scales — pixel-level (close-up cracks), object-level (structural elements such as columns or cornices), and structure-level (whole buildings) — with image resolutions ranging from 168 × 300 to 4,600 × 3,070 pixels and COCO-style instance mask annotations.

The Bai et al. (2021) corpus serves a defined complementary role to the DFM training set, and its inclusion directly addresses a structural limitation identified in preliminary ablation experiments. Although higher-severity classes (S4–S5) were classified with reasonable confidence in initial training runs, consistent with the wider literature, the lower-severity grades S1 and S2 proved substantially harder to discriminate. These grades, characterised by thin cracks and incipient spalling, are precisely the damage manifestations for which the Bai et al. (2021) corpus provides the densest and most finely annotated examples. Initialising the backbone from a checkpoint pretrained on this corpus therefore supplies the feature extractor with crack- and spalling-sensitive filters before any fine-tuning on the DFM data, compensating for the relative scarcity of low-grade examples in the target set.

This pairing justifies the joint use of a small, heterogeneous, EMS-aligned target dataset with a larger, finely annotated source dataset. While the DFM corpus provides domain-specific and operationally grounded supervision across the full severity range, the Bai et al. (2021) dataset anchors the representation of the damage cues that are most difficult to learn from limited examples alone.

Figure 3. Sample images from the DFM test set illustrating the range of damage manifestations addressed by the MSM model. Rows correspond to increasing EMS-98 severity, from grade 1/2 (top row) to grade 5 (bottom row); columns illustrate the heterogeneity of acquisition geometries characteristic of opportunistic ground-survey photography. Images obtained from the INGV DFM.

3.1.1 Annotation strategy

All annotations for model training were produced in CVAT.ai and exported in a format compatible with YOLO segmentation. To align the annotation geometry with the morphology of the damage manifestations across EMS-98 severity classes, a mixed-geometry protocol was adopted. For grades S1–S3, dominated by linear and sub-linear features such as thin cracks and spalling, oriented (rotated) bounding boxes were used. Axis-aligned boxes were deemed unsuitable for these classes because diagonal crack patterns would otherwise be enclosed within rectangles containing a substantial proportion of irrelevant background. For grades S4–S5, which encompass partial collapses and rubble piles whose spatial extents are irregular and topologically complex, damage was delineated directly through polygonal masks. Given that the final learning objective is instance segmentation, the oriented bounding boxes produced for classes S1–S3 were subsequently converted into polygonal masks prior to training, ensuring a consistent label format across all five severity classes.

Figure 4. Examples of the mixed-geometry annotation protocol. Oriented bounding boxes are used for linear damage features (grades S1–S3), while polygonal masks are used for spatially irregular damage (grades S4–S5).

3.2 Transfer learning strategy

Model training for the MSM model relies on transfer learning to address two coupled constraints established in Section 3.1: the limited size of the DFM corpus relative to typical deep-learning requirements, and the documented difficulty of discriminating intermediate EMS-98 grades on heterogeneous imagery (Calantropio et al., 2021; Song et al., 2020). The Bai et al. (2021) corpus, introduced above as a multi-scale source domain, is used to pretrain backbone filters for fine-scale damage cues (cracks, spalling) before fine-tuning the head on the EMS-98 five-class target, a strategy well established for damage-detection tasks in low-data regimes (Bouchard et al., 2022; Yang et al., 2021; Zhuang et al., 2020).

The MSM model is built on the Ultralytics YOLOv11 instance-segmentation architecture (yolo11x-seg; approximately 62 M parameters), selected for its single-stage design, which combines competitive accuracy with the inference speed required for operational deployment (Zou et al., 2024). The single-stage formulation is well aligned with the post-event-only operational regime of macroseismic surveys, which do not assume the availability of pre-event reference imagery.

The architecture follows the standard backbone–neck–head composition: the backbone provides hierarchical feature extraction; the neck (Feature Pyramid Network and Path Aggregation Network) fuses multi-scale features to support the detection of damage manifestations across the wide size range present in macroseismic photography (from millimetric cracks to entire collapsed structures); and the head produces the task-specific outputs for the five-class severity formulation (Xing et al., 2022; Zou et al., 2024).

Transfer learning was implemented in two phases. In the first phase, the YOLOv11x-seg architecture was trained on the Bai et al. (2021) crack-and-spalling dataset, replicating the original benchmark configuration of the study and producing a checkpoint (hereafter bai_best.pt) whose backbone filters are sensitised to fine façade-damage cues, characteristic of EMS-98 grades S1 and S2. In the second phase, the checkpoint was used to initialise the MSM model backbone, while the task was reformulated as a five-class severity classification problem aligned with the EMS-98 by replacing the head with a five-output configuration.

3.3 Training and validation

3.3.1 Training configuration

The MSM model was fine-tuned on the DFM-derived dataset using a staged-unfreezing schedule of three steps. A common training configuration was applied across all three stages. The principal training hyperparameters are summarised in Table 3.

Parameter	Value	Role in training
Data (`data`)	`DFM_YAML`	Dataset configuration (paths, splits, classes).
Image size (`imgsz`)	1280	Input resolution for training/validation (px).
Batch size (`batch`)	2	Images per optimisation step.
Optimizer (`optimizer`)	SGD	Parameter update algorithm.
Final LR fraction (`lrf`)	0.01	Terminal learning-rate fraction at scheduler end.
Momentum (`momentum`)	0.9	SGD term accumulating gradients to stabilise updates.
Weight decay	1×10^{-4}	L2 regularisation strength.
Cosine LR schedule	True	Enables cosine annealing of the learning rate.
Rotation (`degrees`)	5	Max rotation augmentation (deg).
Translation (`translate`)	0.05	Max shift augmentation (fraction of image size).
Scaling (`scale`)	0.5	Scale augmentation range (zoom factor).
Shear (`shear`)	3	Max shear augmentation (deg).
Perspective (`perspective`)	0.0005	Perspective jitter magnitude (kept minimal for linear crack features).
Overlapping masks	True	Allows/handles overlapping instance masks.
Mask ratio (`mask_ratio`)	4	Internal mask downsampling ratio.
Box loss weight (`box`)	7.5	Weight of bounding-box regression loss.
Classification loss weight	0.5	Weight of class prediction loss (five classes).
DFL loss weight (`dfl`)	0.7	Weight of Distribution Focal Loss for box refinement.
Early stopping (`patience`)	50	Epochs without improvement before stopping.
Validation (`val`)	True	Computes validation metrics during training.

Across the three stages, the number of frozen backbone layers is progressively reduced and the initial learning rate is correspondingly lowered, allowing the network to adapt progressively to the representations of the DFM dataset. The staged schedule was designed as follows:

- Stage A (freeze = 12, 25 epochs, $lr_0 = 2.5 \times 10^{-3}$) realigns the inherited Bai checkpoint to DFM imagery characteristics with a deliberately short, high-rate schedule.
- Stage B (freeze = 8, 100 epochs, $lr_0 = 1.0 \times 10^{-3}$) is the primary optimisation phase, allowing mid-level extractors to specialise.
- Stage C (freeze = 0, 150 epochs, $lr_0 = 5 \times 10^{-4}$) performs full fine-tuning at a low learning rate, supporting calibration and finer discrimination among severity grades.

At the end of every epoch the model was evaluated on the held-out validation set, and the best-performing checkpoint was retained for transfer to the next stage. The Stage C best checkpoint constitutes the final MSM model. Standard instance-segmentation metrics were computed separately for bounding-box (B) and mask (M) outputs: precision, recall, mean Average Precision at IoU = 0.5 (mAP@0.5), and mean Average Precision averaged over IoU thresholds from 0.5 to 0.95 (mAP@0.5:0.95). Class-wise metrics across the five EMS-98 grades are reported alongside aggregate scores in Section 4.

3.3.2 Evaluation metrics

Model performance is assessed using four standard classification metrics: accuracy, precision, recall, and F1-score. Each metric is derived from the counts of True Positives (TP), False Positives (FP), False

Negatives (FN), and True Negatives (TN) produced by the model across the test set. Together, they capture both the model's ability to detect damage instances and its tendency to misclassify non-damage regions, and are defined as follows:

$$Accuracy = (TP + TN) / (TP + TN + FP + FN)$$

$$Precision = TP / (TP + FP)$$

$$Recall = TP / (TP + FN)$$

$$F1\text{-}score = 2 \cdot (Precision \cdot Recall) / (Precision + Recall)$$

Operationalisation: building-level aggregation and geodataframe extraction3.4

The trained MSM model produces, for each input image, a set of detected damage instances each characterised by a pixel-level mask, an EMS-98 severity label (S1–S5), and a confidence score. Translating these instance-level predictions into a data product compatible with macroseismic survey practice requires two further processing steps: (i) aggregation of detected instances into a single damage grade per building, and (ii) georeferencing and export of the building-level records as a GIS-interoperable geodataframe.

Building-level aggregation. The aggregation rule adopted by the MSM is consistent with standard EMS-98 macroseismic survey practice, in which a building is assigned an overall grade according to the most severe damage manifestation observed (Grünthal, 1998). Concretely, for each building all damage instances detected across the available photographic views are grouped, and the building-level grade is assigned as the maximum EMS-98 severity among the detected instances exceeding a configurable confidence threshold τ. The threshold parameter exposes a sensitivity / false-positive trade-off that is calibrated against ground-truth grades. The performance metrics for the building-level grade on the held-out DFM test partition are reported in Section 4 alongside the instance-level results.

Geodataframe extraction. Although the MSM model is trained and validated on opportunistic ground-survey imagery (DFM), the surrounding pipeline is conceptually designed to operate on UAV-acquired imagery accompanied by minimal exterior-orientation metadata: GNSS-derived UAV position (latitude, longitude, altitude) and the camera viewing geometry (gimbal- and IMU-derived attitude). Subject to standard UAV-grade GNSS accuracy this metadata is sufficient to project image-space detections onto a horizontal plane and attribute each detection to the nearest building within a predefined Area of Interest (AOI). *T*he resulting building-level records are exported as a GeoPackage (.gpkg) layer in a standardised Coordinate Reference System, with each record carrying:

- a unique building identifier;
- the assigned EMS-98 damage grade (S1–S5);
- the aggregate prediction confidence;
- the supporting UAV exterior-orientation metadata for each contributing image (GNSS latitude, longitude, altitude AGL; gimbal yaw, pitch, roll), as illustrated in Table 9;
- the platform and sensor identifiers (UAV model, focal length, sensor designation);
- the linked source image identifier(s) and acquisition timestamp(s) supporting the assignment.

Additionally, where authoritative building-footprint datasets are available, the projected detections can be intersected with footprint polygons to refine building-level attribution and absorb residual spatial uncertainty introduced by GNSS positioning.

The resulting layer is directly loadable in QGIS or analogous GIS environments and can be overlaid onto authoritative building-footprint datasets to produce per-building EMS-98 grade maps suitable for civil protection agencies and macroseismic surveying. At the present stage of development, this pipeline is implemented and validated on the opportunistic DFM corpus, demonstrating the end-to-end operational viability of the workflow; full deployment testing on live UAV survey acquisitions, alongside model refinement and scalability, is reserved for subsequent work and discussed in Section 5.

4. Results and Discussion

This section reports and discusses the empirical evaluation of the MSM. Section 4.1 presents the development of the instance segmentation model: a preliminary ablation comparing a fully disaggregated 10-class taxonomy with the consolidated 5-class severity scheme adopted in this work (Section 4.1.1), followed by the validation performance of the transfer-learning schedule (Section 4.1.2). Section 4.2 reports

the building-level performance obtained after applying the maximum-grade aggregation rule to the held-out test partition, which is the metric of operational interest for macroseismic surveying. Section 4.3 discusses the implications of these results for deployment within civil-protection and macroseismic-survey workflows. Throughout the discussion, the absolute scores reported should be read against the heterogeneity of the underlying DFM corpus — composed of opportunistic ground-survey photographs exhibiting wide variability in viewpoint, distance, illumination, framing, and occlusion — illustrated in Figure 3.

4.1 Damage segmentation model results

4.1.1 Preliminary ablation: 10-class versus 5-class taxonomy

The EMS-98 distinguishes between masonry (M) and reinforced-concrete (RC) structural typologies, each with five severity grades (M1–M5; RC1–RC5), yielding a fully disaggregated 10-class taxonomy. A preliminary classification ablation was conducted to assess whether this taxonomy could be supported by the available DFM training data, by training a ConvNeXt-base image-level classifier on the same dataset under both taxonomy formulations. The results are reported in Tables 5 and 6.

Table 3: Preliminary image-level classification ablation under the 10-class disaggregated taxonomy (M1–M5, RC1–RC5). Test macro-F1 = 0.294; accuracy = 0.198; weighted-F1 = 0.196; $n = 109$.

Class	Precision	Recall	F1-score	Support
M1	1.000	0.200	0.333	5
M2	0.000	0.000	0.000	22
M3	0.000	0.000	0.000	23
M4	0.000	0.000	0.000	17
M5	0.778	0.875	0.824	8
RC1	0.014	1.000	0.028	4
RC2	1.000	0.400	0.571	6
RC3	0.667	0.500	0.571	10
RC4	0.000	0.000	0.000	7
RC5	0.500	0.800	0.615	7
Macro avg	0.396	0.378	0.294	109
Weighted avg	0.251	0.198	0.196	109

Table 4: Preliminary image-level classification ablation under the consolidated 5-class severity-only taxonomy (S1–S5). Test macro-F1 = 0.600; accuracy = 0.646; weighted-F1 = 0.638; $n = 109$.

Class	Precision	Recall	F1-score	Support
S1	0.500	0.167	0.250	9
S2	0.600	0.667	0.632	28
S3	0.606	0.645	0.625	33
S4	0.667	0.632	0.649	24
S5	0.846	0.846	0.846	15
Macro avg	0.644	0.591	0.600	109
Weighted avg	0.642	0.646	0.638	109

Tables 5 and 6. Preliminary image-level classification ablation under the disaggregated 10-class taxonomy (top) and the consolidated 5-class severity-only taxonomy (bottom). Test set: n = 109.

The contrast between the two formulations is unambiguous. Under the 10-class taxonomy, the intermediate damage grades are essentially unrecoverable: M2, M3, and M4 all collapse to F1 = 0.000, and only the extreme grades (M5, RC5) reach usable F1. Under the consolidated 5-class taxonomy, all intermediate grades become learnable — S2, S3, and S4 reach F1 ≈ 0.62–0.65, and S5 reaches F1 = 0.846 — with macro-F1 doubling (0.294 → 0.600) and overall accuracy more than tripling (0.198 → 0.646). The S1 class remains the weakest (F1 = 0.250), reflecting its very small support and inherent ambiguity with the no-damage baseline, but is no longer collapsed to zero.

Two factors drive the gap. First, consolidating M and RC reduces the effective fragmentation of the training distribution, thus increasing per-class support, which is decisive in a low-data regime. Second, the typology

distinction (M vs. RC) often relies on contextual cues such as surface texture and structural element geometry, which are not always recoverable from a single ground photograph framed on a damage feature. The result is consistent with the body of literature that has reduced finer EMS-98 schemes to coarser ones for the same reason: Calantropio et al. (2021) reduced the EMS-98 scheme from five to two classes in their UAV-based pipeline, citing the insufficient thematic accuracy of intermediate EMS-98 classes from remotely sensed imagery; Song et al. (2020) report analogous findings on aerial damage-mapping tasks. The 5-class severity-only formulation was therefore adopted for the instance segmentation model, with the trade-off that the building-typology information is lost — a limitation discussed in Section 4.3.

4.1.2 Training of the instance segmentation model

The MSM segmentation model was fine-tuned under the three-stage transfer-learning schedule introduced in Section 3.3, initialised from a YOLOv11x-seg checkpoint pretrained on the crack-and-spalling dataset (Bai et al., 2021). The schedule progressively unfreezes the backbone (freeze = 12 → 8 → 0) and lowers the initial learning rate (lr_0 = 2.5 × 10^{-3} → 1.0 × 10^{-3} → 5.0 × 10^{-4}). Validation was performed at the end of every epoch, and the best checkpoint of each stage, selected on validation mAP@0.5, was carried forward to the next.

Stage A exhibits a signature of a successful transfer-learning warm-up (Figure 5). Training box-loss decreases across the 25 epochs, while validation box-loss drops sharply in the first 3–5 epochs as the head realigns to the DFM imagery, reaches a minimum around epoch 9, and exhibits a mild upward drift towards the end of the schedule. The train–val gap remains narrow throughout, indicating that the inherited backbone is already informative for the target domain and that the short, high-rate schedule recovers an initialisation suitable for the subsequent optimisation phase.

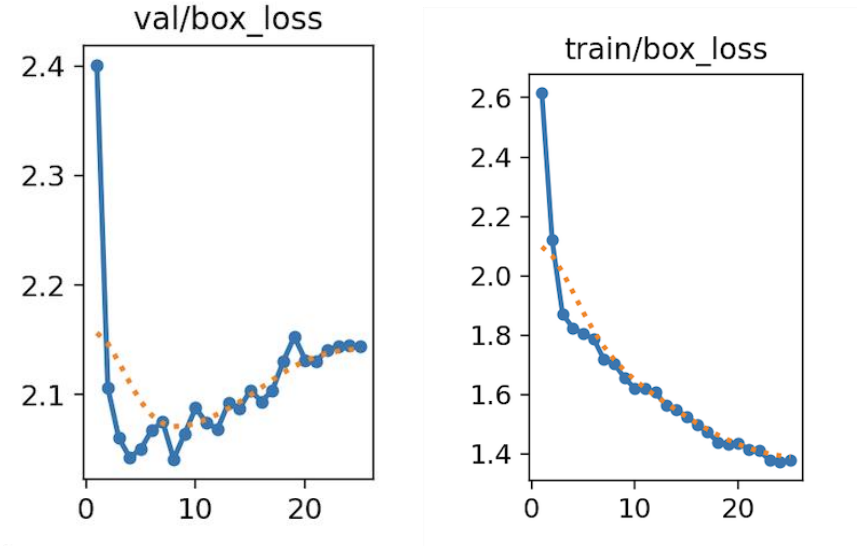

The progression of validation mAP@0.5 across the three stages (Figure 6) is consistent with the staged-unfreezing rationale and with the low-data regime characterising the DFM corpus. In Stage A, mAP rises from near-zero and plateaus around 0.13–0.15, indicating fast adaptation of the head while the inherited backbone is preserved. In Stage B, partial backbone unfreezing yields a further sustained gain, with mAP stabilising around 0.20–0.21. In Stage C, the curve becomes flat and noisy, oscillating around the level reached at the end of Stage B with no substantive further improvement, consistent with the model having converged to the capacity supported by the limited training set. A parallel pattern is observable on the per-stage training and validation losses: training losses continue to decrease across all stages, while validation losses begin to drift upward from early in Stage B onwards. The Stage C best checkpoint, attained at epoch 24, was retained as the final MSM model.

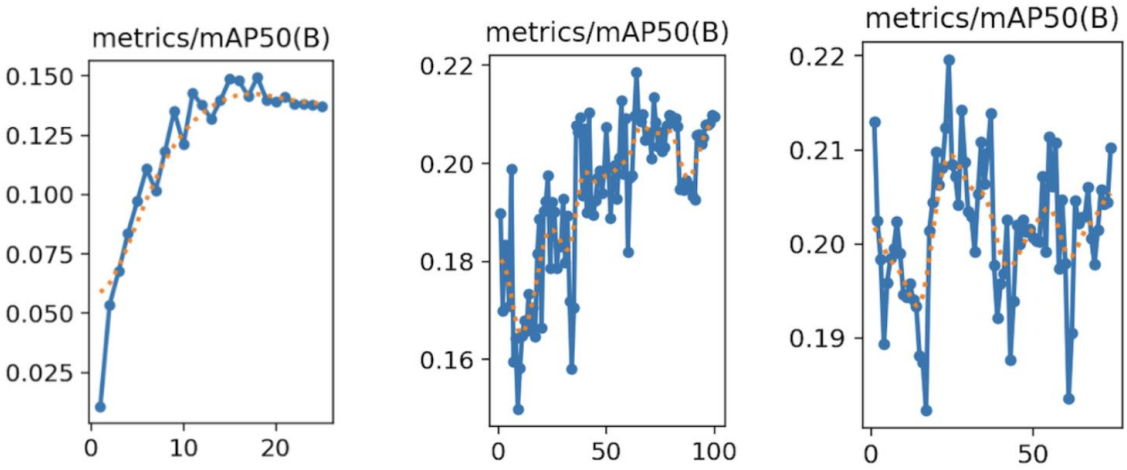

Figure 6. Validation bounding-box mAP@0.5 across the three transfer-learning stages: Stage A (left, 25 epochs), Stage B (centre, 100 epochs), Stage C (right, 150 epochs).

The validation scores reported in Table 7 are moderate in absolute terms and are systematically lower for masks than for boxes (F1 = 0.437 for boxes vs. 0.336 for masks; mAP@0.5 = 0.288 vs. 0.191). The gap is most pronounced for fine, linear damage manifestations such as cracks, where pixel-level mask agreement is intrinsically difficult to achieve, and for large damage entities (rubble piles, partial collapses), whose irregular boundaries make tight delineation particularly sensitive to acquisition noise and to the rough nature of the polygonal annotation strategy adopted on the DFM imagery (Section 3.1.1). The gap between mAP@0.5 (0.288 / 0.191) and mAP@0.5:0.95 (0.075 / 0.057) further indicates that the model achieves usable detection at a permissive IoU criterion but struggles to maintain high-IoU mask precision. This is the operating regime in which the building-level aggregation step, evaluated in Section 4.2, becomes informative: a maximum-grade aggregation rule is comparatively insensitive to mask boundary precision and more sensitive to the recall of severity-discriminative detections.

Table 5: Evaluation metrics for the damage segmentation model using the Ultralytics evaluation protocol. Metrics are reported separately for bounding boxes (B) and instance masks (M).

Metric	Boxes (B)	Masks (M)
Precision	0.468	0.357
Recall	0.409	0.318
F1 score	0.437	0.336
mAP@0.5	0.288	0.191
mAP@0.5:0.95	0.075	0.057

Table 7. Validation metrics for the final MSM model under the Ultralytics evaluation protocol, reported separately for bounding-box (B) and instance-mask (M) outputs.

The absolute validation scores must be interpreted against three properties of the training corpus discussed in Section 3.1. First, the DFM is small (587 training images) for a high-capacity instance-segmentation network; the present model therefore stands as a prototype demonstrating that a five-grade EMS-98 detector can be trained under tight data and compute constraints, rather than as an attempt to saturate the achievable performance on this task. Second, the corpus is opportunistic: photographs were collected by expert macroseismic surveyors during past Italian earthquakes for documentation purposes, with no controlled acquisition geometry, no enforced stand-off distance, and no consistent framing. Third, intra-class visual variability is consequently very wide which broadens the support of each EMS-98 class in feature space and the overlap between adjacent classes. The transfer-learning strategy from the Bai et al. (2021) corpus mitigates these constraints but does not remove them. The scores in Table 7 should therefore be read as a feasibility demonstration on a deliberately challenging dataset, providing a sound prototype on which subsequent UAV-targeted refinements can build on.

SHOW SOMEWHERE THE ms TAKEN TO DETECT EACH IMAGE TO SHOW THAT IT CAN DETECTS VERY FAST!

Beyond the model itself, an explicit objective of this study is to conceptualise and demonstrate the end-to-end MSM pipeline for macroseismic surveys. Read against this objective, the moderate instance-level validation scores are consistent with the dataset characteristics discussed above; the more informative evidence for operational viability comes from the building-level aggregation evaluated in Section 4.2, where instance-level confusion is partially absorbed by the maximum-grade aggregation rule that mirrors EMS-98 surveyor practice.

Figure 7 shows representative MSM model predictions on test imagery, organised in the same row-wise order as the sample images in Figure 3 (increasing EMS-98 severity from top to bottom) and exhibiting visual diversity across viewing geometries (left to right). Each overlay reports the instance-segmentation mask, bounding box, severity-class label, and confidence score produced by the model, illustrating its behaviour both across varying damage degrees and across the acquisition scales typical of the DFM corpus.

Figure 7. MSM model predictions on the same set of test images shown in Figure 3, organised in the same row-wise order of increasing EMS-98 severity (S2 to S5) and column-wise variation of viewing geometry.

4.2 Damage aggregation: building-level performance on the test set

While the metrics in Section 4.1.2 quantify the segmentation model's performance at the instance level, the unit of analysis in macroseismic survey practice is the building, not the individual damage manifestation (Grünthal, 1998; Tertulliani et al., 2025). The MSM pipeline therefore aggregates instance-level predictions into a single per-building EMS-98 grade by applying the maximum-grade aggregation rule introduced in Section 3.4: each building is assigned the highest EMS-98 severity among the detected instances exceeding the confidence threshold τ.

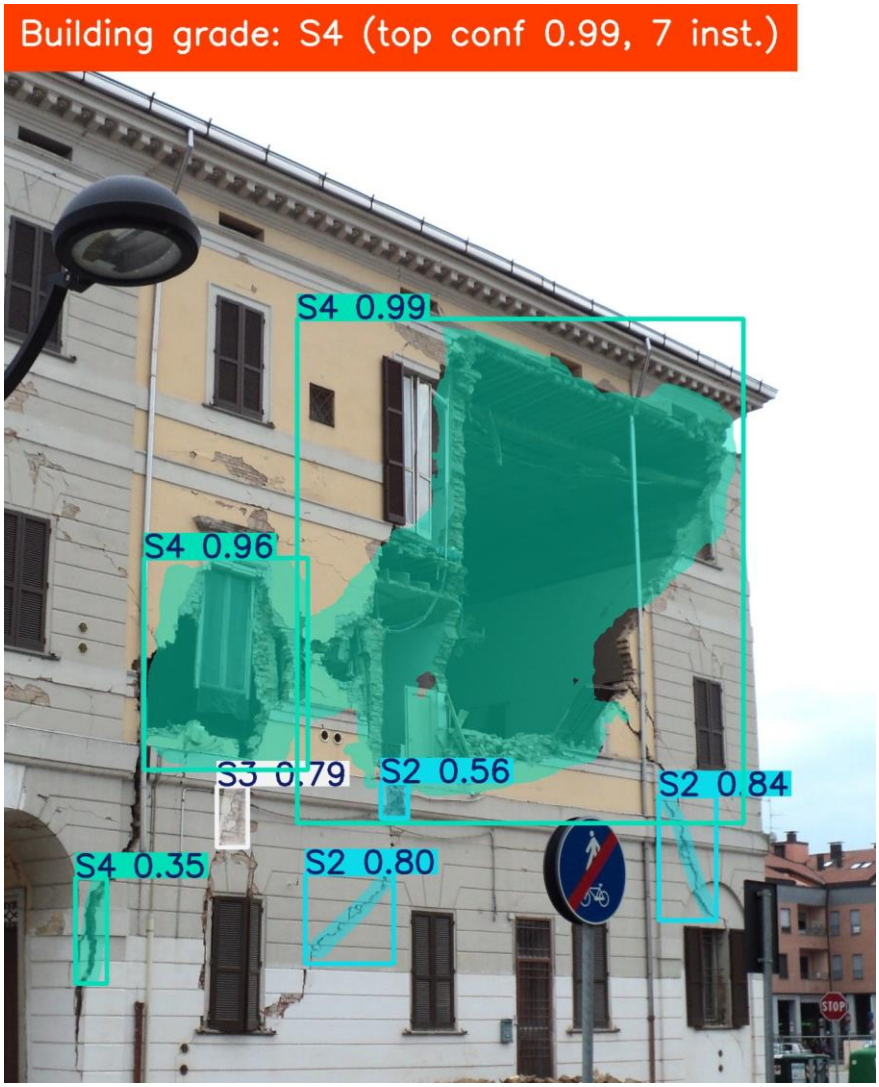

This subsection evaluates the result of that aggregation on the held-out DFM test partition (n = 109). Buildings for which the model produced no detection given that no damage manifestation was present are assigned to the residual class S0. Each test image is treated as a distinct building in the present evaluation; multi-view aggregation per physical structure is examined in Section 4.3. The test set spans the full EMS-98 severity range, with the imbalanced distribution inherited from the DFM corpus reported in Table 2. The resulting confusion matrix between ground-truth and aggregated predicted grades is shown in Figure 8, and the corresponding per-class metrics are reported in Table 8.

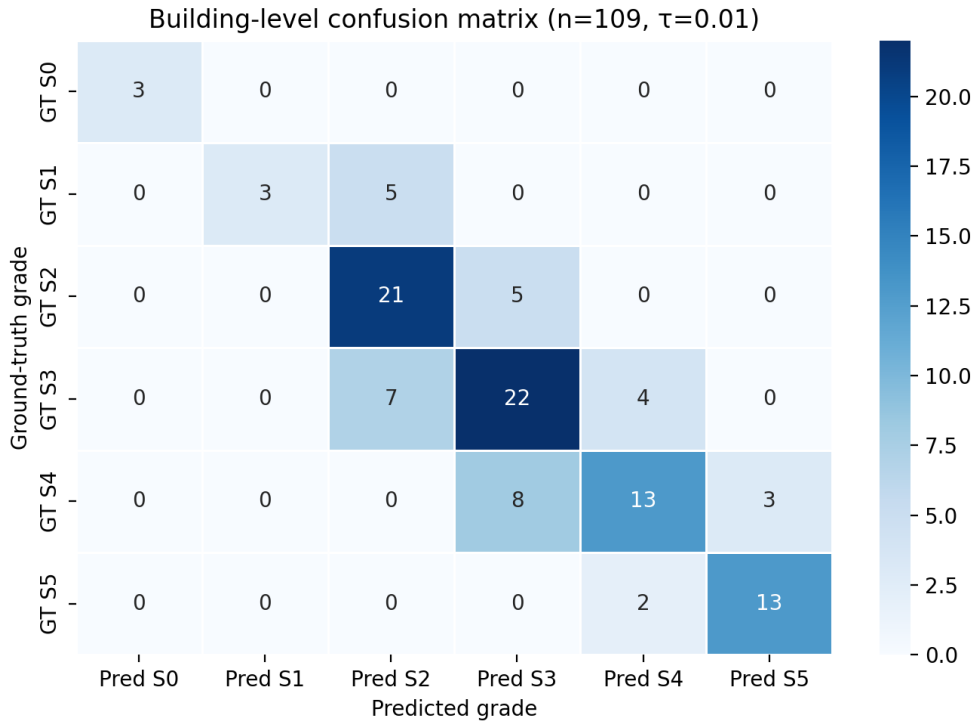

Figure 8. Building-level confusion matrix between ground-truth and aggregated predicted EMS-98 grades on the DFM test partition (n = 109, τ = 0.01). Rows correspond to ground-truth grades and columns to predicted grades.

Class	Support	TP	FN	FP	Precision	Recall	F1-score
S1	*8*	*3*	*5*	*0*	*1.000*	*0.375*	*0.545*
S2	*26*	*21*	*5*	*12*	*0.636*	*0.808*	*0.712*
S3	*33*	*22*	*11*	*13*	*0.629*	*0.667*	*0.647*
S4	*24*	*13*	*11*	*6*	*0.684*	*0.542*	*0.605*
S5	*15*	*13*	*2*	*3*	*0.812*	*0.867*	*0.839*
Macro avg	***106***	—	—	—	***0.752***	***0.652***	***0.670***

Table 8. Per-class building-level performance for severity grades S1–S5: support, true positives (TP), false negatives (FN), false positives (FP), precision, recall, and F1-score. Macro-averaged scores reported in the bottom row. Computed from the confusion matrix in Figure 8.

Three observations summarise the building-level results. First, exact-match accuracy across all five grades is 67.9% (72/106 buildings with ground-truth S1–S5). Even this figure understates the operational utility of the pipeline, because EMS-98 grades are ordinal and adjacent classes are visually similar by construction — a difficulty acknowledged across the literature on damage classification over ordinal scales (Calantropio et al., 2021; Song et al., 2020). Within macroseismic survey practice itself, differences of one intensity degree between independent expert assessments of the same locality are documented and treated as resolvable through revision rather than as outright errors (Tertulliani et al., 2025). When evaluated against

the inherent ordinal structure of the EMS-98 scale, no building in the S1–S5 portion of the test set is misclassified by more than one grade relative to its ground-truth assignment.

Second, the model is symmetric across the under- and over-grading axes (17 buildings under-graded and 17 over-graded; 16.0% each on the S1–S5 portion of the test set), and no building with ground-truth in S1–S5 was assigned the residual class S0. The latter property is the operationally most desirable failure mode: the model never assigns the no-damage label to a damaged structure, and in particular never misses an S4 or S5 building.

Third, per-class performance is asymmetric across the severity range. The S5 class (destruction / collapse) is the best-recalled (0.867) and second-highest in precision (0.812), confirming that severe damage carries a distinctive visual signature that the network captures reliably. The S1 class achieves the highest precision (1.000) but the lowest recall (0.375): the three S1 predictions returned by the model are all correct, but five of the eight S1 buildings are systematically over-graded into S2. This is consistent with the very limited S1 representation in the training corpus (Section 3.1) and with the inherent ambiguity of S1 — "negligible to slight damage" in EMS-98 terminology — which abuts the no-damage baseline and provides few discriminative cues at façade resolution.

Reframed in operational terms, the model correctly flags 90.3% (65/72) of buildings carrying actionable damage (ground truth S3 or higher) as S3 or higher, while raising false alarms (over-grading buildings with ground truth S1–S2 to S3 or higher) on only 14.7% (5/34) of the lower-grade buildings. The combination of a high actionable-damage recall, the absence of severe-damage misses to the residual no-damage class, and the bounded one-grade extent of the residual misclassifications suggests that the MSM, even in its present prototype form, is a promising basis for a triage-oriented macroseismic damage-mapping tool — with further refinement expected from training on a larger and more representative dataset that better matches the intended UAV-deployment regime. The current model should therefore be read not as a replacement for human surveyor verification, but as a first-pass instrument whose role in civil protection workflows is consistent with the supporting position envisaged for AI-based assessment more broadly (Jozi et al., 2024; Kallas and Napolitano, 2025).

4.3 Operability and deployment considerations

The end-to-end character of the MSM pipeline — from instance-level damage detection, through building-level aggregation, to GIS-interoperable export — constitutes the principal contribution of this work. The pipeline provides the full conceptual structure of an automated macroseismic-survey workflow, while the MSM model embedded within it represents a first prototype: trained on a deliberately challenging opportunistic dataset, the model demonstrates the technical viability of EMS-98-aligned segmentation under the constraints encountered in macroseismic-survey practice, and provides a baseline on which UAV-targeted refinements can be developed in subsequent work. The metrics reported in Sections 4.1 and 4.2 should therefore be read not in isolation but in light of how the pipeline is intended to be deployed within macroseismic-survey workflows. This subsection discusses (i) the implications of the dataset characteristics for the absolute performance scores reported above, (ii) the role of the building-level aggregation rule in

absorbing instance-level confusion, (iii) the suitability of the geodataframe output for integration into civil-protection workflows, (iv) the limitations exposed by the present evaluation, and (v) the positioning of the MSM relative to comparable approaches.

4.3.1 Dataset heterogeneity and the interpretation of absolute scores

The performance scores reported in Sections 4.1.2 and 4.2 must be read against the background of the training and test data. The DFM corpus is, by design, an opportunistic dataset: photographs were collected by expert macroseismic surveyors during past Italian earthquakes for the purpose of documenting damage in the field, not for the purpose of training a deep-learning model. The corpus exhibits substantial variability along several axes — viewpoint and stand-off distance, illumination and shadowing, focal length, framing (whole façade vs. single damage feature), background clutter, and partial occlusion. This heterogeneity is intrinsic to the operational acquisition conditions of macroseismic surveys and was deliberately retained as a defining property of the dataset (Section 3.1).

From a learning standpoint, this heterogeneity defines an upper bound on the absolute scores attainable on this corpus. Two consequences should be made explicit. First, intra-class visual variability on the DFM is wider than on benchmark damage-detection datasets curated under controlled-acquisition or single-event conditions, broadening the support of each EMS-98 class in feature space and, consequently, the overlap between adjacent classes. The validation-set instance-segmentation scores (mAP@0.5 ≈ 0.288 / 0.191 for boxes / masks) and the building-level exact-match accuracy of 67.9% should therefore be interpreted in this light. Second, the limited size of the DFM corpus (866 images in total, with 109 in the test partition) places this work in a low-data regime in which both the variance of the reported metrics and the sensitivity to specific test-partition draws are non-trivial. The transfer-learning strategy from the Bai et al. (2021) corpus mitigates this constraint but does not remove it. These results are interpreted as a demonstration of feasibility rather than as the upper bound of attainable performance: deliberate matching of the training data to the deployment regime — UAV-acquired imagery with controlled flight geometry, discussed in Section 4.3.4 — is expected to substantially improve the scores reported here, and is the principal direction for future work.

4.3.2 The role of building-level aggregation

The maximum-grade aggregation rule adopted in the MSM pipeline is consistent with EMS-98 surveyor practice, in which a building is assigned the grade corresponding to the most severe damage manifestation observed (Grünthal, 1998). For an automated pipeline operating on noisy instance-level outputs, this rule has two operationally relevant properties. First, it is robust to instance-level under-confidence: provided that at least one instance of the correct severity is detected with sufficient confidence, the building is correctly graded, irrespective of how many lower-grade detections accompany it. This is reflected in the per-image data, where buildings with multiple detections (typically several instances per building) generally

include at least one detection in the correct severity tier even when other detections are misclassified into adjacent grades.

Second, the rule absorbs a substantial part of the instance-level confusion that drives the relatively modest segmentation-only metrics reported in Table 7. The increase from instance-level mAP@0.5 (0.288 for boxes) to building-level exact-match accuracy (67.9%) reflects this absorption directly: many instance-level misclassifications cancel out at the building level once the maximum-grade rule is applied, because the operationally relevant signal is the presence of at least one correctly identified severe instance, not the precise grading of every detection.

4.3.3 Geodataframe output and integration with macroseismic workflows

As mentioned in Section 2.3, operational emergency-management workflows require structured, geospatially explicit data products that can be loaded directly into the GIS environments routinely used in civil protection operations rooms. The terminal output of the MSM pipeline is designed to meet this requirement: a building-level GeoPackage layer in which each record carries the assigned EMS-98 grade, an aggregate confidence score, the per-grade detection counts, and the linked supporting imagery. This output distinguishes the MSM pipeline from the bulk of the deep-learning damage-detection literature, in which model outputs typically remain at image-level or visualisation-level. As Kallas and Napolitano (2025) document, the absence of a structured data product to support decision-making is a recurring barrier to the operational uptake of AI-based assessment tools, irrespective of the accuracy attained on benchmark datasets; Jozi et al. (2024) reach a parallel conclusion, identifying the absence of end-to-end automation from acquisition to georeferenced damage records as a principal obstacle to rapid post-disaster deployment. The MSM responds to both observations directly, by structuring its outputs around the unit of analysis used by macroseismic surveyors.

Two design choices follow from this orientation. First, the per-record retention of the supporting imagery and per-grade detection counts preserves the audit trail necessary for human verification: a surveyor can inspect the underlying imagery and segmentation masks for any building flagged at a given grade, allowing the automated assessment to be reviewed and, where necessary, corrected. This is critical in macroseismic practice, where the assigned intensity is ultimately the responsibility of an expert team (Tertulliani et al., 2011, 2025) and where automated outputs must remain inspectable. Second, the GeoPackage format was selected for its open, single-file portability and its native support in QGIS, the GIS environment most commonly used by Italian civil-protection actors and macroseismic-survey teams.

Image	*Timestamp (UTC)*	*Lat (°)*	*Lon (°)*	*Alt AGL (m)*	*Gimbal yaw (°)*	*Gimbal pitch (°)*	*Gimbal roll (°)*	*Focal (mm)*	*Sensor*	*Aggregate Prediction*
1464.jpg	*2009-04-12 09:14:00*	*42.35044*	*13.39963*	*38*	*192.0*	*-30*	*0.1*	*12.29*	*L2D-20c 4/3″ CMOS*	*S1*
1629.jpg	*2009-04-12 09:14:47*	*42.35030*	*13.39995*	*42*	*228.0*	*-35*	*-0.2*	*12.29*	*L2D-20c 4/3″ CMOS*	*S2*
1743.jpg	*2009-04-12 09:16:39*	*42.34978*	*13.40003*	*47*	*300.0*	*-32*	*0.1*	*12.29*	*L2D-20c 4/3″ CMOS*	*S3*
1953.jpg	*2009-04-12 09:17:38*	*42.34959*	*13.39975*	*40*	*336.0*	*-38*	*-0.1*	*12.29*	*L2D-20c 4/3″ CMOS*	*S3*
2573.jpg	*2009-04-12 09:19:43*	*42.34970*	*13.39905*	*50*	*48.0*	*-34*	*0.0*	*12.29*	*L2D-20c 4/3″ CMOS*	*S4*
2975.jpg	*2009-04-12 09:21:41*	*42.35022*	*13.39897*	*44*	*120.0*	*-36*	*0.1*	*12.29*	*L2D-20c 4/3″ CMOS*	*S5*
3224.jpg	*2009-04-12 09:22:42*	*42.35041*	*13.39925*	*39*	*156.0*	*-31*	*0.0*	*12.29*	*L2D-20c 4/3″ CMOS*	*S5*

Operationally, the rows of Table 9 are transformed into the GeoPackage layer through four steps that constitute the geodataframe-extraction component of the MSM pipeline. First, for each row the GNSS triplet (lat, lon, alt AGL) together with the gimbal yaw and pitch defines the camera's optical-axis ray, which is intersected with the ground plane to yield a projected ground point in WGS-84 — for example, the first row of Table 9 projects to a ground point approximately 66 m south-south-west of the drone position. Second, each projected point is attributed to a building: where authoritative cadastral footprints are available (as in most Italian historic centres), the point is matched to the polygon it falls within or, failing direct intersection, to the nearest polygon centroid within a configurable spatial buffer; in the absence of footprints, the projected point is exported as a point feature and the building identity is resolved later. Third, where multiple table rows project to the same building — the normal case under a façade-orbiting flight pattern — the per-image predictions in the rightmost column are aggregated by the maximum-grade rule already adopted at the instance level (Section 3.4), so that each building carries the most severe EMS-98 grade observed across the available views. Fourth, the resulting per-building records are written to a single GeoPackage (.gpkg) file in which the geometry column stores the matched footprint (or projected point) and the attribute schema retains the source image identifier(s), the supporting metadata fields from Table 9, the per-grade detection counts, the aggregate confidence score, and the assigned EMS-98 grade. The output is a standard OGC GeoPackage layer that can be loaded directly into QGIS or any GIS environment supporting the format, where it can be styled by grade, queried, intersected with other

layers, or used as input to downstream analyses. Figure 9 illustrates the rendered result for a representative AOI.

Figure 9. Illustrative GeoPackage output of the MSM pipeline rendered in QGIS. Building footprints are coloured by aggregated EMS-98 damage grade; supporting imagery and instance-level detection overlays remain linked from each record, preserving the audit trail required by macroseismic-survey practice. The pie chart summarises the distribution of buildings across damage grades for the depicted area of interest.

4.4. Limitations and constraints on operational deployment

Class imbalance and the S1 grade. The DFM corpus is severely imbalanced toward intermediate and severe grades, with only 44 S1 images in total. The systematic over-grading of S1 buildings into S2 on the test set (5 of 8 S1 buildings) is a direct consequence of this imbalance and of the inherent ambiguity of S1, which sits at the boundary with the undamaged baseline. In an operational context, this is partially mitigated by the fact that S1 buildings — by EMS-98 definition associated with negligible structural impact — are the lowest-priority class for emergency response, and over-grading them into S2 introduces no actionable consequence. Calibrated S1 detection would nonetheless be required for full macroseismic intensity assessment, in which the share of buildings at each grade enters the intensity computation directly (Grünthal, 1998); achieving this would require targeted augmentation of the S1 sample, the introduction of an explicit undamaged class, or a hybrid scheme in which no-damage outcomes are inferred rather than predicted.

Loss of EMS-98 typology and building-vulnerability information. The 5-class severity-only formulation adopted on the basis of the ablation reported in Section 4.1.1 discards the masonry / reinforced-concrete typology distinction that is intrinsic to the EMS-98. This is a methodologically motivated trade-off — the data do not support a 10-class formulation in the present state — but it is also a substantive limitation: a complete EMS-98 macroseismic assessment requires the share of buildings of each vulnerability class at each damage grade. Beyond construction-material typology, a further attribute that the present pipeline does not estimate is the EMS-98 vulnerability class of each building, which jointly with the damage grade determines the contribution to the intensity computation. Estimating vulnerability from imagery alone is non-trivial even for trained ground surveyors and is typically inferred from contextual information (age of the building stock, construction era, design code). Future work in which typology and vulnerability are recovered through complementary classifiers operating on the building footprint and on auxiliary cadastral / building-stock layers, or through targeted expansion of the training corpus, is a natural extension of the present pipeline.

Domain shift between training and deployment imagery. The DFM corpus consists of opportunistic ground-survey photographs taken by human surveyors with handheld cameras: short stand-off distances, eye-level oblique perspectives, and operator-selected framing of damage features. UAV-derived imagery in the intended deployment scenario will differ systematically — longer stand-off distances, more uniform near-orthogonal façade perspectives, lower per-pixel resolution at equal image size, and an absence of

operator framing, so that whole-façade views with peripheral damage cues become the norm. The model has not been exposed to imagery with these acquisition characteristics during training, and the metrics in Sections 4.1.2 and 4.2 should therefore be interpreted as in-domain estimates rather than as a direct estimate of UAV-deployment performance. Future work should prioritise the assembly of a UAV-derived, EMS-98-annotated validation set, even of modest size, to characterise this domain shift quantitatively.

5. Conclusion

This study has presented the Macroseismic Survey Mapper (MSM), a conceptualized end-to-end pipeline for the automated generation of damage maps aligned to the EMS-98 grading. The contribution is twofold. At the workflow level, the MSM is a conceptual and operational pipeline that links the four stages of a post-event survey — disaster onset, UAV imagery acquisition, automated damage detection and building-level aggregation, and GIS-interoperable export — into a single, traceable chain which terminal output is a building-level GeoPackage layer. At the model level, the MSM model is a YOLOv11x-seg instance-segmentation network trained, under a staged transfer-learning schedule from the Bai et al. (2021) crack-and-spalling corpus, on the heterogeneous and deliberately challenging INGV Database Fotografico Macrosismico (DFM).

The evaluation reported in Section 4 supports two complementary conclusions. First, on a corpus that imposes a structural ceiling on attainable scores — small size, opportunistic acquisition, wide intra-class variability — the MSM model attains an exact-match building-level accuracy of 67.9% and recovers actionable damage (ground-truth S3 or higher) in 90.3% of cases with a false-alarm rate of 14.7% on lower-grade buildings. Second, beyond the absolute scores, the principal demonstration of the study is the operational viability of the full pipeline end-to-end on real macroseismic imagery, with the building-level aggregation rule absorbing a substantial share of the instance-level confusion that limits the segmentation-only metrics. The MSM model is therefore best understood as a feasibility prototype rather than as a saturated solution: it establishes that an EMS-98-aligned, building-level macroseismic-survey output can be produced from a tractable, low-resource training pipeline, and it provides a baseline on which UAV-targeted refinements can build.

Three directions for future work follow directly from the analyses in Section 4.3. First, the assembly of a UAV-derived, EMS-98-annotated validation set — even of modest size — would allow the domain shift between handheld ground-survey imagery and operational UAV acquisitions to be characterised quantitatively, and would provide the training material required to lift the absolute scores of the instance segmentation model. Second, targeted augmentation of the S1 sample, the introduction of an explicit undamaged class trained on uncompromised façade imagery, or a hybrid inference scheme for no-damage outcomes, would address the residual ambiguity at the boundary between S1 and the no-damage baseline. Third, the packaged software tool delivered alongside this manuscript is intended to support the operational

deployment of the pipeline on live UAV surveys, and a controlled pilot deployment is a next step to validate the workflow under operational conditions.

Taken together, these results position the MSM as a feasible foundation for an operational, GIS-interoperable, AI-supported macroseismic-survey workflow, with a clearly identified path from the present prototype to UAV-grade deployment.

References

Ainscoe, E.A., et al. (2025). Earthquake damage mapped more comprehensively and accurately by radar satellites than optical imagery. Communications Earth & Environment, 6(1), 631.

Al Shafian, S., Hu, D. (2024). Integrating Machine Learning and Remote Sensing in Disaster Management: A Decadal Review. Buildings, 14(8), 2344.

Bai, Y., Sezen, H., Yilmaz, A. (2021). End-to-end Deep Learning Methods for Automated Damage Detection in Extreme Events at Various Scales. ICPR 2020, pp. 6640–6647.

Bouchard, I., et al. (2022). On Transfer Learning for Building Damage Assessment from Satellite Imagery. Remote Sensing, 14(11), 2532.

Brunner, D., Lemoine, G., Bruzzone, L. (2010). Earthquake Damage Assessment of Buildings Using VHR Optical and SAR Imagery. IEEE Transactions on Geoscience and Remote Sensing, 48(5), 2403–2420.

Calantropio, A., et al. (2021). Deep Learning for Automatic Building Damage Assessment. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, V-1-2021, 113–120.

Cheng, C.-S., et al. (2024). A framework to enhance disaster debris estimation with AI and aerial photogrammetry. International Journal of Disaster Risk Reduction.

Duarte, D., et al. (2017). Towards a more efficient detection of earthquake-induced façade damages using oblique UAV imagery. ISPRS Archives.

Fernandez Galarreta, J., Kerle, N., Gerke, M. (2015). UAV-based urban structural damage assessment using object-based image analysis and semantic reasoning. Natural Hazards and Earth System Sciences, 15(6), 1087–1101.

Gerke, M., Kerle, N. (2011). Automatic Structural Seismic Damage Assessment with Airborne Oblique Pictometry Imagery. Photogrammetric Engineering & Remote Sensing, 77(9), 885–898.

Grünthal, G. (Ed.) (1998). European Macroseismic Scale 1998. Cahiers du Centre Européen de Géodynamique et de Séismologie, Luxembourg.

Jozi, S., et al. (2024). Rapid post-disaster assessment of residential buildings using Unmanned Aerial Vehicles.

Kakooei, M., Baleghi, Y. (2017). Fusion of satellite, aircraft, and UAV data for automatic disaster damage assessment. International Journal of Remote Sensing, 38(8), 2511–2534.

Kallas, J., Napolitano, R. (2025). Image-To-Insight: A novel workflow for converting post-disaster imagery of historic masonry structures into actionable data. International Journal of Disaster Risk Reduction, 120, 105358.

Kerle, N., Nex, F., Gerke, M., Duarte, D., Vetrivel, A. (2019). UAV-Based Structural Damage Mapping: A Review. ISPRS International Journal of Geo-Information, 9(1), 14.

Khan, A., Gupta, S., Gupta, S.K. (2022). Emerging UAV technology for disaster detection, mitigation, response, and preparedness. Journal of Field Robotics, 39(6), 905–955.

Kim, B., Cho, S. (2020). Automated Multiple Concrete Damage Detection Using Instance Segmentation Deep Learning Model. Applied Sciences, 10(22), 8008.

Kim, J., Leite, F. (2026). A unified framework for automated damage assessment in post-disaster built environments using LiDAR.

Matin, S.S., Pradhan, B. (2022). Challenges and limitations of earthquake-induced building damage mapping techniques using remote sensing images. Geocarto International, 37(21), 6186–6212.

Rossi, A., et al. (2019). The 2016–2017 earthquake sequence in Central Italy: macroseismic survey and damage scenario through the EMS-98 intensity assessment. Bulletin of Earthquake Engineering, 17(5), 2407–2431.

Song, D., et al. (2020). Integration of super-pixel segmentation and deep-learning methods for evaluating earthquake-damaged buildings using single-phase remote sensing imagery. International Journal of Remote Sensing, 41(3), 1040–1066.

Tavakol Sadrabadi, S., et al. (2024). Conceptual Design of a Wildfire Emergency Response System Empowered by Swarms of UAVs.

Tertulliani, A., et al. (2011). An application of EMS-98 in a medium-sized city: the case of L'Aquila (Central Italy) after the April 6, 2009 Mw 6.3 earthquake. Bulletin of Earthquake Engineering, 9(1), 67–80.

Tertulliani, A., et al. (2025). A comprehensive integrated macroseismic dataset from multiple earthquake studies. Earth System Science Data, 17(8), 4063–4077.

Wiguna, S., et al. (2024). Evaluation of Deep Learning Models for Building Damage Mapping in Emergency Response Settings. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 17, 5651–5667.

Xing, L., Fan, X., Dong, Y., Xiong, Z., Xing, L., Yang, Y., Bai, H., Zhou, C. (2022). Multi-UAV cooperative system for search and rescue based on YOLOv5. International Journal of Disaster Risk Reduction, 76, 102972.

Yang, W., Zhang, X., Luo, P. (2021). Transferability of Convolutional Neural Network Models for Identifying Damaged Buildings Due to Earthquake. Remote Sensing, 13(3), 504.

Zhuang, F., et al. (2020). A Comprehensive Survey on Transfer Learning. arXiv:1911.02685.

Zou, R., et al. (2024). An Improved Instance Segmentation Method for Fast Assessment of Damaged Buildings Based on Post-Earthquake UAV Images. Sensors, 24(13), 4371.