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

Giovanni Galli[a,*], Marco Dubbini[a], Filippo Bernardini[b], Luca Arcoraci[b]

[a] *Department of History and Cultures (DiSCi), University of Bologna, Piazza San Giovanni in Monte, 2, Bologna, 40124, Italy*
[b] *Istituto Nazionale di Geofisica e Vulcanologia (INGV), Sezione di Bologna, Via Carlo Berti Pichat, 6/2, Bologna, 40127, Italy*

Abstract

Post-earthquake macroseismic surveying often relies on ground-based visual inspections that are slow, costly, and difficult to scale in the immediate aftermath of a seismic event. Deep-learning damage detectors have advanced substantially in recent years, yet their outputs are rarely translated into operational deployment tools that yield a georeferenced dataframe of buildings aligned with the European Macroseismic Scale 1998 (EMS-98). This study presents the Macroseismic Survey Mapper (MSM), a prototype end-to-end pipeline that links unmanned aerial vehicle imagery acquisition, instance-level damage segmentation across the five EMS-98 severity grades, per-building aggregation under a worst-observed rule, and export to an OGC GeoPackage layer ready for direct ingestion into geographic information systems. On the held-out test set the model reaches 64.6% exact-match accuracy; treated as a binary triage classifier, it recalls 90.3% (65/72) of buildings carrying actionable damage (EMS-98 grade S3 or higher) while over-grading only 13.5% (5/37) of lower-grade buildings, with all but two mis-classifications falling within one grade of the ground truth. The pipeline is exercised on a real UAV deployment over Piedilama (Arquata del Tronto, Central Italy), mapping the surveyed block at 3.78±0.32 s per image. The MSM establishes a deployable end-to-end workflow; future work targets UAV-native training data to quantify operational accuracy at scale.

Keywords: damage assessment, UAV imagery, deep learning, EMS-98, macroseismic survey, GIS pipeline

*Corresponding author
Email addresses: giovanni.galli7@studio.unibo.it (Giovanni Galli), marco.dubbini@unibo.it (Marco Dubbini), filippo.bernardini@ingv.it (Filippo Bernardini), luca.arcoraci@ingv.it (Luca Arcoraci)

1. Introduction

1.1. Problem context

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 [1, 2]. 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 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 [3, 4].

Despite its operational importance, post-earthquake damage assessment remains predominantly reliant on ground-based visual inspections conducted by trained personnel [5, 6]. 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 [5, 6]. 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 while coverage remains sparse elsewhere [7].

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 [8, 9]. 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 [10, 11, 12, 13].

In parallel, computer vision techniques have demonstrated strong performance in damage detection, classification, and segmentation tasks, with modern single-stage architectures reaching levels of accuracy and efficiency compatible with operational deployment [14, 15, 16]. 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 [2, 17]. Consequently, a persistent gap remains between image-derived damage detection and its translation into actionable information employable by the emergency actors involved in damage assessment workflows.

This limitation is particularly evident in macroseismic assessment contexts [18, 19, 5]. 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 [20, 21]. Moreover, most methods do not produce outputs in formats readily interoperable within Geographic Information Systems (GIS), limiting their integration into operational decision-making [6, 17].

1.2. Aim and contributions

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 on building's façades, followed by building-level damage aggregation, and (iv) export of a geodataframe (GeoPackage) of the damage in the surveyed locality. 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. Consistent with this distinction, the MSM is presented here as a prototype that prioritises speed, ease of use, and interoperability over saturated detection accuracy: the contribution lies in demonstrating an operable damage-mapping workflow that can be scaled up to larger, UAV-native datasets, with the current segmentation model providing a moderately accurate interpretive capacity. The pipeline is presented end-to-end and tested in a deployment scenario over the village of Piedilama (Arquata del Tronto, Central Italy); the case study is reported in Section 4.3.

1.3. Paper structure

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 and wider emergency workflows. Section 3 describes the dataset, model architecture, training strategy, and pipeline design, including the

geospatial attribution chain that converts UAV-derived image-level detections into a building-level damage map. Section 4 presents and discusses the experimental results, including the case study in Central Italy. Section 5 concludes with implications for operational deployment and directions for future work.

2. Background

2.1. Remote sensing 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 [22, 1, 2]. 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 [8, 9]. 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 [23].

Thus, 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 [11, 12]. A growing body of work has demonstrated the value of directly processing oblique UAV imagery for façade damage assessment. Fernandez et al. [13] 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. [20] 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. [6] 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. [24] 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 6070% 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 [25, 17, 26], 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 [27, 28, 15, 2]. 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 [29] and has been demonstrated to perform well on post-event imagery alone, without recourse to pre-event reference data [15]. 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 [14, 30]. 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 [31, 32]. 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 [30] 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 are limited in 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 [33, 34, 31].

2.3. GeoAI Tools for operational deployment

A persistent limitation across the damage detection literature is the disconnect between model outputs and their operational use. Emergency management workflows generally require explicit geospatial 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 [6, 32, 17]. 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 [3, 5, 7, 21]; 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 [20, 35, 36]. Indeed, distinguishing adjacent damage degrees in the intermediate range (grades 1, 2 and 3) is a challenge even for professional ground surveyors [4]. 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 information products. Kallas and Napolitano [17] 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. [6] raise a parallel concern, identifying the absence of end-to-end automation 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 [7, 5, 4]. Notably, the scale's original specification explicitly anticipated a basis for computerised macroseismic intensity evaluation [3], 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 [17] 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.

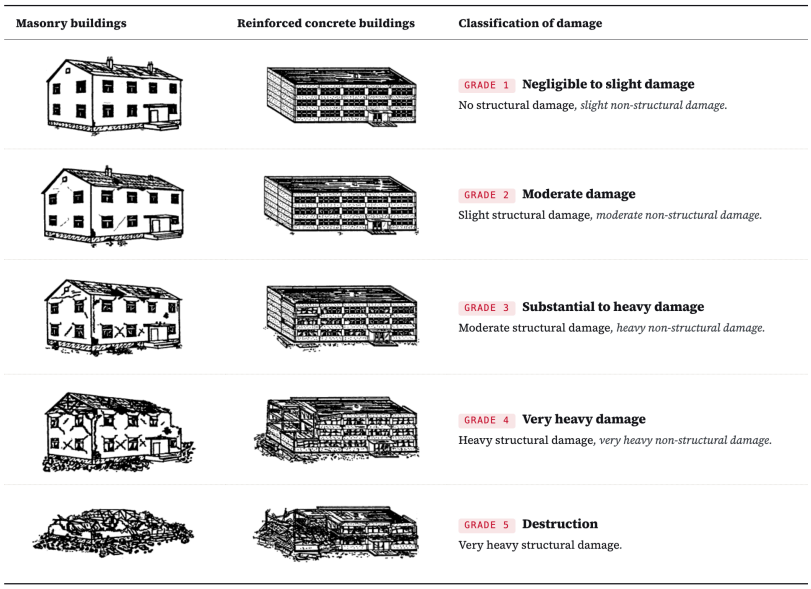

Figure 1: The EMS-98 damage grading scheme used as the reference taxonomy in this work. The scale defines five severity grades and distinguishes between masonry (M) and reinforced-concrete (RC) structural typologies. Adapted from [3].

2.4. Research gap 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) previous models on damage detection and classification aligned to the EMS-98 fail to classify across all the five damage degrees and 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 per-building grades nor an auditable evidence trail.

In response to these gaps, this study develops the Macroseismic Survey Mapper (MSM) introduced in Section 1. The end-to-end MSM pipeline is summarised in Figure 2 and is organised into four operational stages. **Stage 1 (Event)** marks the onset of the seismic disaster at time T_0 and triggers the workflow. **Stage 2 (Drone survey)** consists of a UAV mission, typically flown one to two weeks post-event over the affected area along façade-oriented orbits; its output is the set of acquired

images $\{I_k\}$ together with the camera pose recorded in the EXIF/XMP metadata (latitude, longitude, altitude, yaw, pitch, roll). **Stage 3 (MSM model)** is the analytical core of the pipeline: the image set is processed by the instance-segmentation network, which detects damage manifestations on each façade, classifies each instance across the five EMS-98 severity grades, and aggregates the instance-level predictions into a single damage grade per building. **Stage 4 (GeoPackage export)** fuses the per-building damage grades with the camera-pose metadata to attribute each detection to a building footprint and to write the resulting layer as an OGC GeoPackage (`damage.gpkg`), the deliverable consumed downstream by GIS-based emergency-management workflows.

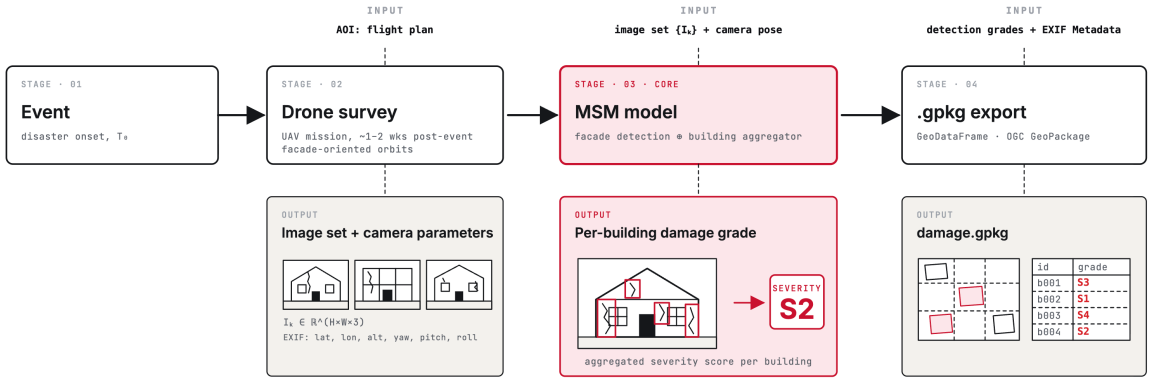

Figure 2: The MSM end-to-end pipeline: (i) the seismic event triggers the workflow; (ii) a UAV mission acquires oblique façade imagery with EXIF/XMP metadata; (iii) the MSM model segments, classifies (EMS-98 S1–S5), and aggregates damage per building; (iv) detections are exported as a GIS-interoperable GeoPackage.

The pipeline is exercised end-to-end on a UAV deployment over the village of Piedilama (Arquata del Tronto), at the epicentre of the 2016–2017 Central Italy seismic sequence and assigned an EMS-98 intensity of XI ("devastating") in the official macroseismic survey of the event [4, 7]. The village was left largely as it stood in the immediate aftermath, making it a rare site at which the pipeline's outputs can be compared against a documented macroseismic baseline. The case study is reported in Section 4.3 and demonstrates that the workflow yields a building-level damage map at an average of 3.78 ± 0.32 s per image.

3. Materials and Methods

This study proposes a five-step methodology to develop the MSM. The methodology comprises: (1) compilation of a heterogeneous training corpus from the INGV Database Fotografico Macrosismico (DFM)[1]; (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 pre-trained on the Bai et al. [14] 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 geo-referenced dataframe through UAV GNSS and camera-orientation metadata.

The remainder of this section details each step of the pipeline: the training corpus and its annotation are described in Sections 3.1 and 3.2; the transfer-learning strategy and training schedule in Sections 3.3 and 3.4; and the operationalisation chain — from instance detections to the building-level GeoPackage — in Section 3.5.

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 machine 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 4.236.29 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 3.3.

Although the EMS-98 distinguishes between masonry (M) and reinforced-concrete (RC) structural typologies, an initial fully disaggregated ten-class formulation produced fragmented sample

[1] Available at https://dfm.ingv.it (last accessed: May 28, 2026).

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 DFM by epicentral area and event.

sizes that compromised the learning of intermediate grades. The two construction-type taxonomies were therefore consolidated into a unified five-class severity scheme (S1S5), a simplification consistent with strategies adopted in comparable studies [35, 36]. An ablation study formally comparing the disaggregated 10-class formulation with the unified 5-class scheme 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 after merging masonry (M) and reinforced-concrete (RC) classes into unified damage levels (S1–S5).

Figure 3: Sample images from the DFM test set illustrating the range of damage manifestations addressed by the MSM tool. Rows correspond to increasing EMS-98 severity, from grade 1/2 (top row) to grade 5 (bottom row), while columns reflect the heterogeneity of acquisition geometries from close-range views to mid-range façade shots and wider building-scale frames. Images obtained from the INGV Database Fotografico Macrosismico (DFM).

To support transfer learning, training was initialised from a checkpoint derived from the dataset curated by Bai et al. [14], 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 Œ 300 to 4,600 Œ 3,070 pixels and COCO-style instance mask annotations. The Bai et al. [14] 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 (S4S5) 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. [14] 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 filters sensitive to low severity grades 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 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. [14] dataset anchors the representation of the damage cues that are most difficult to learn from limited examples alone.

3.2. Annotation protocol

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 S1S3, 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 S4S5, 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 S1S3 were subsequently converted into polygonal masks prior to training, ensuring a consistent label format across all five severity classes.

Figure 4: Mixed-geometry annotation protocol in CVAT.ai: oriented bounding boxes for the linear/sub-linear features of grades S1–S3, polygonal masks for the irregular extents of S4–S5, with all S1–S3 boxes converted to masks prior to training.

3.3. 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 [20, 36]. The Bai et al. [14] 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 [31, 34, 33].

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 [15]. The single-stage formulation is well aligned with the post-event operational regime of macroseismic surveys, which do not assume the availability of pre-event reference imagery. The architecture follows the standard backboneneckhead 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 of the target detections; and the head produces the task-specific outputs for the five-class severity formulation [29, 15].

Transfer learning was implemented in two phases. In the first phase, the YOLOv11x-seg architecture was trained on the Bai et al. [14] crack-and-spalling dataset, replicating the original benchmark configuration of the study and producing a checkpoint (hereafter bai_best.pt) whose backbone fil-

ters 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.

3.4. Training and validation

3.4.1. 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**
Image size (`imgsz`)	1280	Input resolution for training/validation (px).
Batch size (`batch`)	8	Images per optimisation step (explicit, not autobatch).
Mixed precision (`amp`)	True	Enables automatic mixed-precision training.
Optimizer (`optimizer`)	SGD	Parameter update algorithm.
Momentum (`momentum`)	0.9	SGD term accumulating gradients to stabilise updates.
Weight decay (`weight_decay`)	1×10^{-4}	L2 regularisation strength.
Cosine LR schedule (`cos_lr`)	True	Enables cosine annealing of the learning rate.
Final LR fraction (`lrf`)	0.01	Terminal learning-rate fraction at scheduler end.
Overlapping masks (`overlap_mask`)	True	Allows/handles overlapping instance masks.
Mask ratio (`mask_ratio`)	4	Internal mask downsampling ratio.

Table 3: Common training hyperparameters used across all transfer-learning phases.

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 = 2.5×10^{-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 = 1.0×10^{-3}) is the primary optimisation phase, allowing mid-level extractors to specialise.
- Stage C (freeze = 0, 150 epochs, lr_0 = 5×10^{-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.

3.4.2. 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:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{2}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{3}$$

$$F\_1\text{-score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{4}$$

3.5. Operationalisation

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 (S1S5), and a confidence score. Translating these instance-level predictions into a data product compatible with macroseismic survey practice requires three further processing steps: (i) image-level aggregation of detected instances into a single damage grade per photograph, (ii) spatial attribution of each photograph to a building via UAV exterior-orientation metadata, and (iii) building-level aggregation across photographs that may depict the same structure, followed by export of the building-level records as a geodataframe.

3.5.1. Image-level aggregation

The first aggregation step collapses the set of detected damage instances of a facade into a single per-building EMS-98 grade. Following standard EMS-98 practice, where a structure takes the grade

of its most severe observed damage [3], the MSM applies a worst-observed rule: among instances with confidence at least τ, the image grade is the highest EMS-98 severity. Images with no valid instance are assigned the residual no-damage class S0. The building-level performance metrics for this rule on the held-out DFM test partition are reported in Section 4.2 alongside the instance-level results. The output of this step is a per-image record carrying the predicted grade, the confidence of the top detection, and the per-class instance counts (n_S1 n_S5).

3.5.2. Spatial attribution

Although the MSM model is trained and validated on opportunistic ground-survey imagery (DFM), the surrounding pipeline is designed to operate on UAV-acquired imagery accompanied by minimal exterior-orientation metadata: GNSS-derived UAV position (latitude, longitude, altitude AGL) and the camera viewing geometry (gimbal yaw pitch, and roll). Subject to standard UAV-grade GNSS accuracy, this metadata is sufficient to project each image-space detection onto a horizontal plane and to attribute the photograph to a building. This process takes the following three steps:

Camera point. For each photograph, the GNSS latitude/longitude triplet defines the UAV position at the time of acquisition, exported as a point feature in a projected metric Coordinate Reference System.

Projected line of sight. Each photograph defines a 3D camera optical-axis ray, anchored at the UAV's xyz position and pointing in the direction given by the gimbal yaw and pitch. The yaw fixes the direction of the axis in the horizontal plane. The horizontal ground distance from the UAV's planimetric position on the horizontal plane to the point at which the optical axis intersects the horizontal plane is:

$$d_h = \frac{H_{AGL}}{\tan(|\theta_{pitch}|)} \tag{5}$$

with H_{AGL} the drone's above-ground altitude (`RelativeAltitude` in the DJI XMP metadata) and θ_{pitch} the gimbal pitch. The exported geometry is the ground-plane segment of length d_h from the UAV's planimetric position to this hit-point—the 2D footprint of the 3D optical axis, sufficient for footprint intersection and natively renderable in 2D GIS.

Image-to-building attribution. Where authoritative building-footprint datasets are available, each line of sight d_h is intersected with the respective footprint polygon that first intersects it. When a line crosses multiple footprints, the footprint closest to the camera position is selected. Selecting the closest intersected polygon makes the pipeline robust in dense urban fabric, allowing for acquisitions in narrow streets.

3.5.3. Building-level aggregation and GeoPackage export

Given the orientation of the camera towards the facade, an individual building is photographed from several viewpoints. The building-level aggregation step collapses the per-image grades depicting one building into a single record, applying the same worst-observed rule already used at the image level: the building's max_grade is the highest EMS-98 grade attributed to it across all supporting photographs [3]. A mode_grade is also retained as a robustness check against single-image misclassification, together with per-severity instance sums (sum_S1 sum_S5), the number of supporting observations, the mean top confidence, and the list of source image filenames preserved for human inspection. The resulting per-building records are written to a single GeoPackage (.gpkg) layer in a standardised CRS, in which the geometry column stores the matched footprint polygon. Each record carries:

- a unique building identifier inherited from the footprint dataset;
- the assigned EMS-98 damage grade (max_grade, S0S5) and the mode_grade as a robustness attribute;
- the number of supporting observations (n_observations) and the mean top-detection confidence (mean_top_conf);
- per-severity instance sums across all supporting photographs (sum_S1 sum_S5), preserving a per-building damage inventory;
- the filenames of the images carrying the detections (image_ids), so that each individual detection and its EMS-98 classification can be inspected directly.

The output is a standard OGC GeoPackage layer that can be loaded directly into any GIS environment where it can be styled by grade, queried, intersected with other layers, or used as input to downstream analyses. The full pipelineis implemented as a single, configurable Python notebook that operates on a folder of UAV JPGs, an authoritative footprint layer, and the trained model checkpoint; its end-to-end runtime is reported for the Piedilama deployment in Section 4.3.

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 then reports the Piedilama deployment case study, in which the full MSM pipeline is executed end-to-end on UAV imagery. Section 4.4 collects the general limitations of the present prototype. Throughout the discussion, the absolute scores reported should be read against the heterogeneity of the underlying DFM corpus illustrated in Figure 3, and against the framing of the MSM as a prototype whose operational value is established by the end-to-end pipeline rather than by saturated detection metrics.

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 (M1M5; RC1RC5), 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 4 and 5. Under the 10-class taxonomy, the intermediate damage grades are essentially unrecoverable: M2, M3, and M4 all collapse to $F_1 = 0.000$, and only the extreme grades (M5, RC5) reach usable F_1. Under the consolidated 5-class taxonomy, all intermediate grades become learnable—S2, S3, and S4 reach $F_1 \approx 0.62$–0.65, and S5 reaches $F_1 = 0.846$—with macro-F_1 doubling ($0.294 \rightarrow 0.600$) and overall accuracy more than tripling ($0.198 \rightarrow 0.646$). The S1 class remains the weakest ($F_1 = 0.250$), 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 analogous reasons [20, 36]. The 5-class severity-only

formulation was therefore adopted for the instance segmentation model, with the trade-off that information on the building typology is lost a limitation discussed in Section 4.4.

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 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
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

Table 5: 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; $n = 109$.

4.1.2. Training of the instance segmentation model

The MSM instance-segmentation model was fine-tuned under the stage schedule introduced in Section 3.3, initialised from a YOLOv11x-seg checkpoint pretrained on the Bai et al. dataset [14]. The schedule progressively unfreezes the backbone (freeze $= 12 \rightarrow 8 \rightarrow 0$) and lowers the initial learning rate ($\mathrm{lr}_0 = 2.5 \times 10^{-3} \rightarrow 1.0 \times 10^{-3} \rightarrow 5.0 \times 10^{-4}$).

Validation was performed at the end of every epoch, and the best checkpoint of each stage, selected on validation $\mathrm{mAP}_{@0.5}$, 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 35 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 trainval 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.

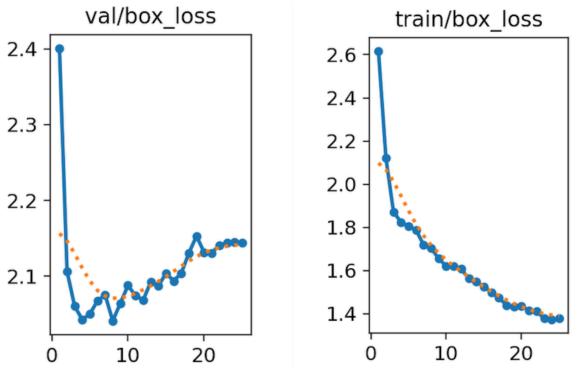

Figure 5: Stage A training and validation box-loss curves over 25 epochs. The training box-loss decreases steadily throughout the schedule, while the validation box-loss drops sharply within the first 3–5 epochs and reaches a minimum around epoch 9 before exhibiting a mild upward drift.

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 0.130.15, 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 0.200.21. 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.

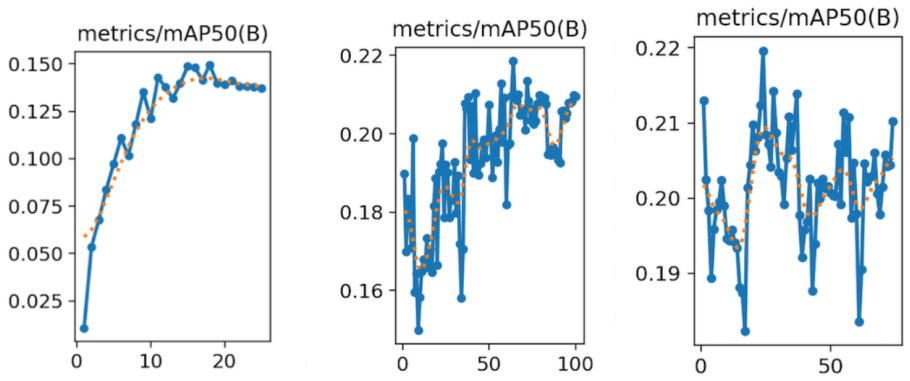

Figure 6: Validation bounding-box mAP@0.5 across the three transfer-learning stages: Stage A (left, 25 epochs), Stage B (centre, 100 epochs), Stage C (right, 150 epochs).

The validation scores reported in Table 6 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. 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.

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 6: 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 read against three properties of the training corpus established in Section 3.1: the DFM is small (587 training images) for a high-capacity instance-segmentation network; it is opportunistic, comprising documentation photographs taken by macro-seismic surveyors with no controlled acquisition geometry; and its intra-class visual variability is

consequently wide, broadening each EMS-98 class in feature space and increasing the overlap between adjacent grades. 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 attempting to saturate achievable performance. The transfer-learning strategy from the Bai et al. [14] corpus mitigates these constraints but does not remove them. The scores in Table 6 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.

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 the following Section and from the deployment case study reported in Section 4.3.

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 [3, 4]. The MSM pipeline therefore aggregates instance-level predictions into a single per-building EMS-98 grade by applying the maximum-grade aggregation rule specified in Section 3.5.1: each building is assigned the highest EMS-98 severity among the detected instances exceeding the confidence threshold τ.

(a) Pianura Emiliana, 2012

(b) Aquilano, 2009

Figure 8: Illustration of the worst-observed aggregation rule on two representative images of the test set. Multiple damage instances of varying EMS-98 severity are detected on each façade by the MSM model; each image is assigned the highest detected grade, with the supporting per-class instance counts retained as auxiliary attributes.

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. 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 9, and the corresponding per-class metrics are reported in Table 7.

Three observations summarise the building-level results. First, exact-match accuracy across all

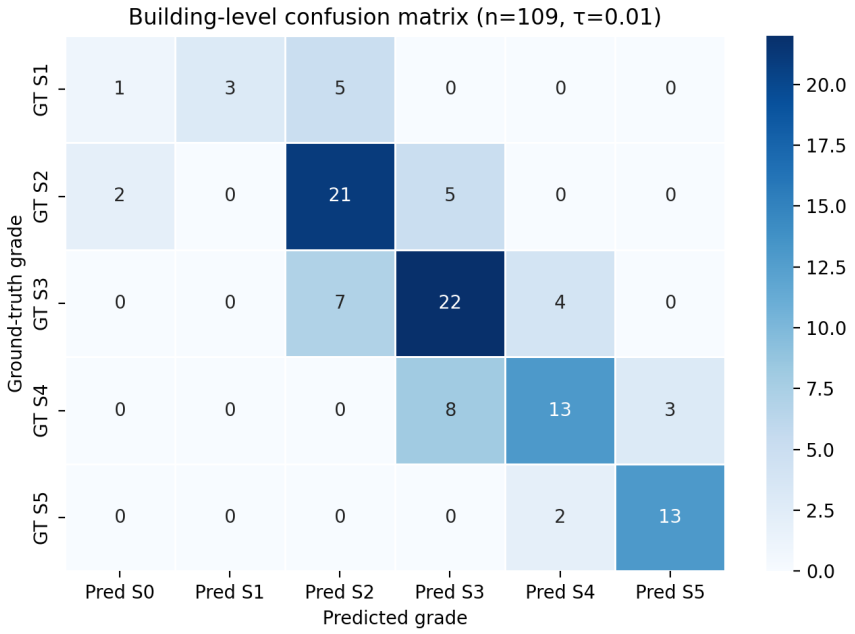

Figure 9: Building-level confusion matrix between ground-truth and aggregated predicted EMS-98 grades on the DFM test partition (n = 109). Rows correspond to ground-truth grades and columns to predicted grades.

five grades is 66.1% (72/109 buildings with ground-truth S1S5). However, this figure understates the pipeline's operational utility. EMS-98 grades are ordinal, and adjacent lower damage grades are visually similar, so a misclassification into a neighbouring grade is far less consequential than the raw accuracy suggests—a difficulty recognised throughout the literature on damage classification over ordinal scales [20, 36]. Indeed, within macroseismic survey practice, one-degree differences between independent expert assessments of the same locality are routinely documented and resolved through revision rather than treated as errors [4]. When evaluated against the inherent ordinal structure of the EMS-98 scale, all but two of the residual mis-classifications fall within one grade of the ground truth; the two exceptions are GT-S2 buildings on which no detection exceeded the confidence threshold and which were consequently assigned the residual no-damage class S0. Second, the model exhibits a moderate asymmetry across the under- and over-grading axes (20 buildings under-graded and 17 over-graded; 18.3% and 15.6% respectively on the $n = 109$ test set). Crucially, all three buildings assigned to the residual no-damage class S0 carry low-severity ground truth (1 GT-S1 and 2 GT-S2): the model never misses an S3, S4 or S5 building to the no-damage class. This is the operationally most desirable failure mode, since missed detections are confined to severity grades that correspond to negligible-to-moderate structural impact and therefore the lowest-priority class for emergency response. Third, per-class performance is asymmetric across the severity range.

Table 7: Per-class building-level performance on the DFM test partition ($n = 109$,). FN counts for S1 and S2 include three buildings (1 GT-S1, 2 GT-S2) assigned to the residual no-damage class S0. Computed from the confusion matrix in Figure 9

Class	Support	TP	FN	FP	Precision	Recall	F1-score
S1	9	3	6	0	1.000	0.333	0.500
S2	28	21	7	12	0.636	0.750	0.689
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	**109**	—	—	—	**0.752**	**0.632**	**0.656**

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 (0.333): the three S1 predictions returned by the model are all correct, but five of the nine S1 buildings are over-graded into S2 and one is assigned the residual class S0. This is consistent with the very limited S1 representation in the training corpus (Section 3.1).

A more operationally meaningful way to assess the model is as a binary triage classifier: does it correctly separate buildings with actionable damage (ground-truth S3 or higher) from those without? On this criterion the model recalls 90.3% (65/72) of actionable-damage buildings, while over-grading only 13.5% (5/37) of lower-grade (S1–S2) buildings to the actionable category. This matters because the operational decision the map supports is not the exact grade but whether a building warrants inspection; for that decision, distinguishing S3-and-above from below is what counts.

Three features support reading the prototype as a viable triage instrument: high recall of actionable damage; no severe-damage buildings misclassified to the residual no-damage class; and bounded errors, with all but two of the misclassified cases falling within one grade of the ground truth. We therefore read the current model not as a replacement for surveyor verification but as a first-pass instrument, consistent with the supporting role envisaged for AI-based assessment more broadly [6, 17]. Further refinement is expected from training on a larger, more representative dataset better matched to the intended UAV-deployment regime.

4.3. Case study: end-to-end deployment over Piedilama (Arquata del Tronto)

The metrics reported in Sections 4.1 and 4.2 quantify what the MSM model can recover on the corpus on which it was trained, but the primary contribution claim of this work is operational:

that an EMS-98-graded, GIS-ready building damage map can be produced end-to-end from a UAV survey within minutes of acquisition. To exercise this claim outside the dataset corpus, the full MSM pipeline was deployed over the village of Piedilama, a small hamlet of the municipality of Arquata del Tronto (Ascoli Piceno, Marche region, Central Italy) that lies at the centre of the area struck by the 20162017 Central Italy seismic sequence and that exhibits a contiguous historic-masonry fabric in varied states of damage. The case study is presented as a successful prototype deployment, not as a quantitative benchmark, since no field-survey ground truth was independently collected for the same buildings on the day of the flight.

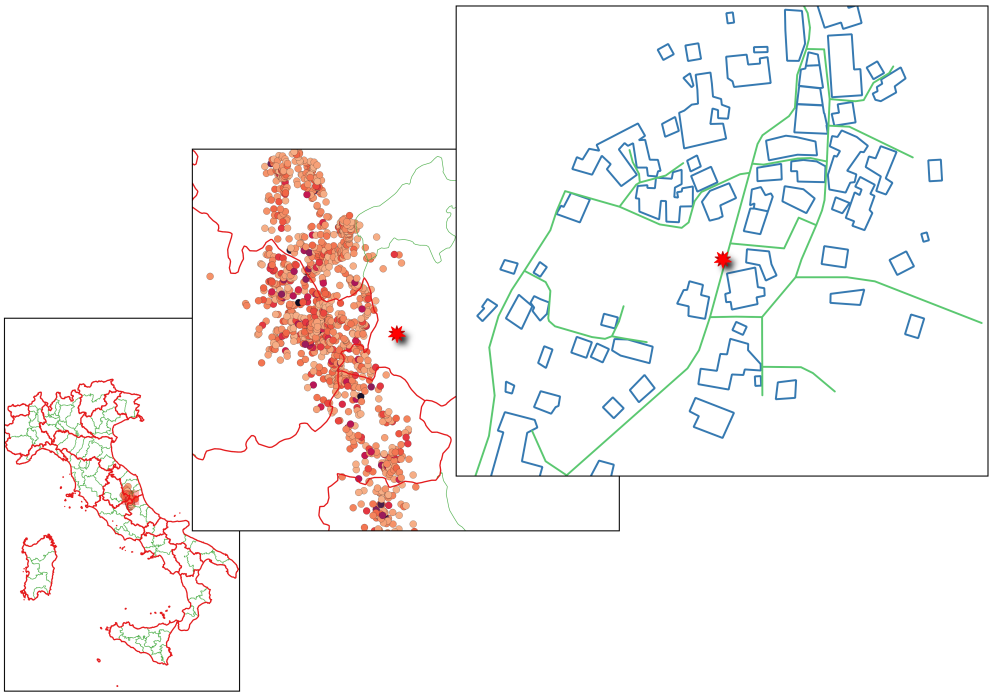

Figure 10: Location of the Piedilama case-study area within the 2016–2017 Central Italy seismic sequence. Left: regional context within Italy. Centre: epicentres of the sequence ($M_w \geq 3$, hypocentral depth ≤ 30 km) shaded by magnitude, with Piedilama marked by the red star. Right: building footprints (blue) and road network (green) of the surveyed block. Seismic Data acquired from [37]

4.3.1. UAV acquisition and study area

A single UAV mission was flown on 22 May 2026 over a contiguous block of the village, covering 18 buildings along the main road. The aircraft was a DJI Mavic 3 (camera model FC3582; focal length 6.72 mm; field of view $\approx 74°$), flown manually at relative altitudes of 14.9–28.1 m AGL with the gimbal pitched downward to oblique angles (typically $-28.7°$ from horizontal) so as to expose

building façades rather than rooftops. The flight produced 161 JPGs with full DJI EXIF and XMP metadata: GNSS triplet (latitude, longitude, ellipsoidal altitude), relative altitude AGL, gimbal yaw/pitch/roll, flight yaw/pitch/roll, camera model, focal length, and field of view. No ground control points were laid and no photogrammetric post-processing was performed: the case study deliberately exercises the data-light branch of the pipeline, in which exterior-orientation metadata alone drives the geospatial attribution. The footprint layer was extracted from the regional CTR for the Marche region, in EPSG:3004 (Monte Mario / Italy zone 2), with the CODICE attribute providing the per-building unique identifier inherited downstream by the GeoPackage output. The location of the deployment area and the perimeter of the surveyed block are shown in Figure 11.

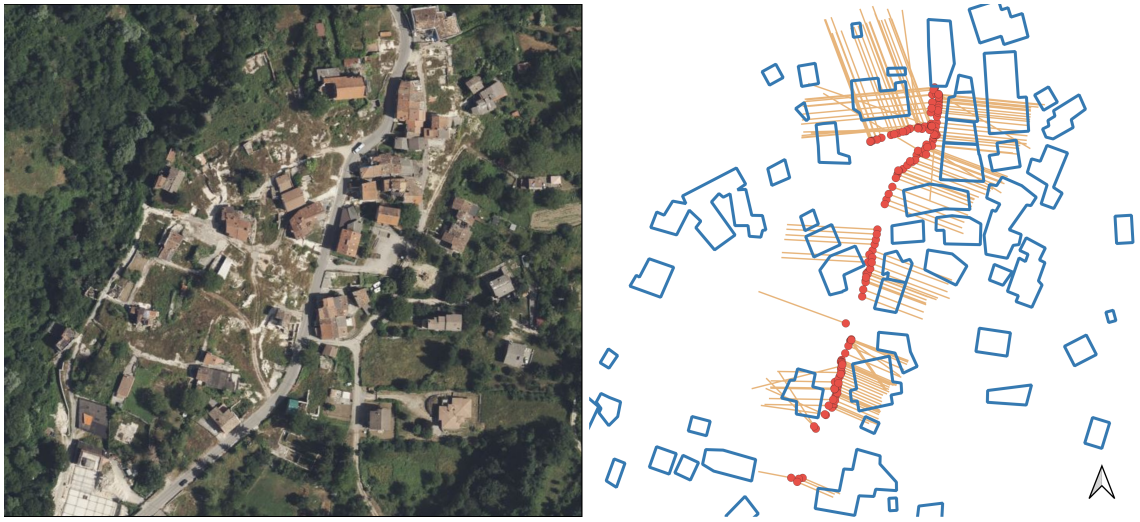

Figure 11: Piedilama UAV deployment area. Left: aerial orthophoto of the village. Right: building footprints (blue) overlaid with UAV camera positions (red points) and the projected lines of sight (orange) reconstructed from the EXIF/XMP exterior-orientation metadata. The lines of sight are clamped to a maximum projection length of 100 m

4.3.2. End-to-end pipeline execution

The 161 images were processed top-to-bottom in the unified MSM notebook, with no parameter tuning beyond the default operational configuration (imgsz = 1280, detector confidence = 0.01, IoU = 0.4, max_det = 300, retina-masks enabled, aggregation threshold $\tau = 0.25$, line-of-sight clamps d_min = 1 m and d_max = 100 m, target CRS = EPSG:3004). Four sub-stages were executed in sequence.

Image-level prediction. The MSM model was loaded once and applied to the images batch. The CPU-only inference pass took 3.78 ś 0.32 s per image, with the entire batch completing in approximately 10 minutes. Each annotated JPG was written to disk with YOLO's instance masks and bounding boxes superimposed, plus a coloured banner stamped with the per-image worst-observed

grade, the top-detection confidence, and the count of valid instances. The per-image grade distribution was S0: 5, S1: 5, S2: 12, S3: 39, S4: 90, S5: 10.

Metadata extraction. EXIF and XMP fields were extracted in bulk via ExifTool, parsed into a pandas dataframe, and exported as drone_metadata.csv. Coordinates were returned in decimal degrees and gimbal/flight angles as plain floats.

Point and line geometry construction. For each image, a camera point was built from the GNSS triplet and a projected line of sight was constructed from the gimbal yaw and pitch using the horizontal-run formula introduced in Section 3.5.2. Line lengths ranged from 14.8 m (a low-altitude near-frontal acquisition of a partially collapsed structure) to 51.3 m (a higher-altitude oblique view of a building at the far edge of the block). Outputs were exported as drone_points.gpkg and drone_lines.gpkg in EPSG:3004.

Image-to-building attribution and building-level aggregation. The 161 lines of sight were intersected with the 70 building footprints, with the rule on closest building handling the cases in which a line crossed two adjacent buildings. 155 of the 161 images were successfully attributed to a footprint (96.3% match rate); the remaining 6 lines of sight either fell outside the coverage or did not intersect any polygon within the clamped projection length. Image-level grades were then joined onto the attributed lines (image_to_building.gpkg) and finally aggregated per building (Damage.gpkg) under the worst-observed rule, with mode grade, per-severity instance sums, mean top-detection confidence, and the comma-separated list of supporting image filenames retained as auxiliary attributes.

Figure 12: Multi-view attribution of the church (`CODICE 337040_3349`): six of the 35 UAV photographsused to classify the building, showing converging lines of sight on different façades. Per-image grades vary across S2–S4; the worst-observed aggregation rule resolves them into a single building-level `max_grade` = S4.

4.3.3. Damage map and attribute table

The terminal output of the pipeline Damage.gpkg is the damage map (Figure 13). The 155 photographs are attributed to 18 distinct buildings (the remaining 52 footprints fall outside the flown block and were left unobserved). The max-grade distribution across the 18 observed buildings is S2: 1, S3: 5, S4: 8, S5: 4, with no building flagged as S1 or S0. The mode grade agrees with the max grade on the buildings with single-grade evidence and diverges only where multiple-view evidence is present, in which case the max grade is the operationally relevant value retained for the damage map. Multi-view evidence is itself a substantive feature of the deployment: several buildings were photographed from two or more positions along the flight track (for instance footprint 337040_3349, observed from 8 images), and in each such case the worst-observed rule recovered the highest grade across the views exactly the behaviour anticipated by the building-level aggregation discussed in Section 4.2.

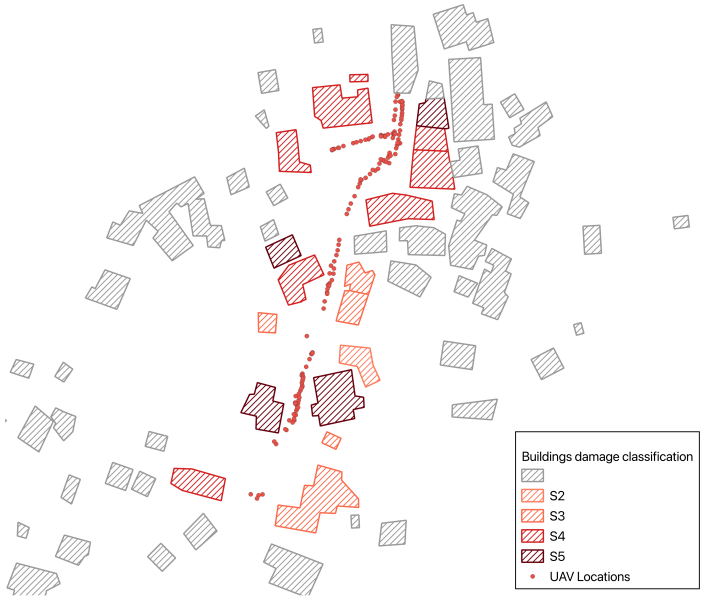

Figure 13: QGIS rendering of the final damage map (Damage.gpkg) over Piedilama, styled by `max_grade` on the EMS-98 colour ramp (S2 yellow → S5 dark red). Drone positions (drone_points.gpkg) and lines of sight (drone_lines.gpkg) are overlaid in light blue to make the supporting evidence visible for each classified building. Unobserved footprints are rendered in light grey for context.

Table 8 reports a representative excerpt of the final building-level attribute table that constitutes the GeoPackage output. The full schema preserves the join key (CODICE), the aggregated grades (max_grade, mode_grade), an evidence count (n_observations), an aggregate confidence (mean_top_conf), per-severity instance sums (sum_S1 sum_S5), and critically for human verifi-

cation the list of supporting source images (image_ids). The complete table for the 18 observed buildings forms the deliverable damage map.

Table 8: Sample of the building-level attribute table exported from the MSM pipeline (subset of 8 rows from the 18 detected buildings in the Piedilama deployment).

CODICE	n_obs	mean_top_conf	max_grade	mode_grade	sum_S1	sum_S2	sum_S3	sum_S4	sum_S5	image_ids
337040_2823	16	0.579	S4	S4	12	11	36	53	0	DJI_0769; DJI_0778
337040_2845	1	0.143	S2	S2	0	2	0	0	0	DJI_0794
337040_2848	9	0.275	S3	S3	1	5	15	0	0	DJI_0682; DJI_0691
337040_2856	27	0.259	S5	S4	29	57	27	25	5	DJI_0583; DJI_0588
337040_2857	4	0.293	S5	S5	0	0	2	9	4	DJI_0929
337040_2860	4	0.443	S3	S3	7	11	7	0	0	DJI_0583
337040_3349	35	0.219	S4	S4	2	127	107	48	0	DJI_0821; DJI_0827; DJI_0829; DJI_0874; ...
337040_29023	11	0.292	S5	S4	0	3	13	17	1	DJI_0002; DJI_0978

4.3.4. Case Study Discussion

Several properties of the deployment are worth making explicit, because they bear directly on the operational claim of the MSM. Speed. The end-to-end runtime for a set of 161 images was on the order of a few minutes on a single CPU, classifying 18 buildings. While the inference took a majority of the time, the geospatial steps added less than a minute combined.

Multi-view robustness. The buildings observed from multiple photographs (n_obs $\geq$ 32 in Table 8) illustrate why the worst-observed aggregation rule is appropriate for UAV deployment: per-image predictions on the same façade frequently spread across two adjacent grades, but the rule recovers the highest of these and assigns it to the building. This design choice is congruent to the convention of the EMS-98. The mode-grade column in the attribute table is retained precisely to make any divergence visible to the human surveyor.

Auditability. Every building record carries the list of source image filenames that contributed to its grade. A surveyor verifying a flagged S4 or S5 building can open the corresponding annotated JPGs, inspect the masks, and confirm or override the automated grade. The audit trail is intrinsic to the data product and aligns the MSM with the principle that the assigned macroseismic intensity is ultimately the responsibility of an expert team [5, 4].

Interoperability. The output is a standard OGC GeoPackage. No specialised software is required to load, query, or style the map: any GIS environment can open the file and the damage map of Figure 13 was produced with a single style-by-attribute operation on max_grade.

Caveats. Two limitations on the strength of the conclusions of the case study should be made explicit. First, no field-survey ground truth was collected on the day of the flight, and the deployment therefore demonstrates operational feasibility rather than per-building classification accuracy in the

deployment regime quantitative validation against an EMS-98-annotated UAV reference set is the principal direction for future work. Second, the MSM model was trained on the opportunistic DFM corpus, whose handheld ground-survey perspective differs systematically from UAV oblique acquisitions; the case study should therefore be read as evidence that the pipeline is operationally viable end-to-end, with the segmentation component still at a moderate-accuracy operating point that will be improved by targeted UAV-native training material in future research.

4.4. Limitations of the study

Four main limitations apply to the present prototype and circumscribe the conclusions that can be drawn from the results reported above.

Dataset heterogeneity and low-data regime. The DFM corpus on which the MSM model was trained 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, and is small in absolute terms (866 images total, 587 in the training partition, 109 in the test partition). This heterogeneity broadens the support of each EMS-98 class in feature space and the overlap between adjacent classes, while the limited size makes the model variance and the sensitivity to specific test-partition draws non-trivial. The transfer-learning strategy from the Bai et al. [14] corpus mitigates these constraints but does not remove them, and the absolute validation-set scores (mAP@0.5 $\approx$ 0.288 / 0.191 for boxes / masks) and the building-level exact-match accuracy of 66.1% should be interpreted in this light. These metrics are best read as a feasibility ceiling on the chosen training corpus rather than as the upper bound of the achievable performance on the task. Targeted assembly of a UAV-native, EMS-98-annotated training set is expected to substantially improve the validation metrics and is the principal direction for future work.

Class imbalance. 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 9 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 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 [3]; 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 is also a substantive limitation: a complete EMS-98 macroseismic assessment requires the share of building typology at each damage grade. Beyond 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 when available). 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 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 Piedilama deployment reported in Section 4.3 exercises the pipeline under this domain shift and shows that the MSM remains operationally usable, but its quantitative effect on per-building classification accuracy in the UAV regime has not been measured. 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 prototype 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 damage geodataframe layer. At the model level, the MSM model is a YOLOv11 instance-segmentation network trained, under a staged transfer-learning schedule from the Bai et al. [14] crack-and-spalling corpus, on the heterogeneous and deliberately

challenging INGV Database Fotografico Macrosismico (DFM).

The evaluation reported in Section 4 supports three 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 66.1% and recovers actionable damage (ground-truth S3 or higher) in 90.3% of cases with a false-alarm rate of 13.5% 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 workflows, with the building-level aggregation rule absorbing a substantial share of the instance-level confusion that limits the segmentation-only metrics. Third, the Piedilama case study reported in Section 4.3 demonstrates that the pipeline runs end-to-end on a UAV deployment, producing a GIS-ready building-level damage map of the surveyed block within minutes of acquisition, and that the worst-observed aggregation rule resolves the multi-view evidence intrinsic to UAV surveys into a single coherent grade per building. The MSM model is therefore best understood as a feasibility prototype rather than as a final product: 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 on.

Three directions for future work follow directly from the analyses in Section 4.3 and Section 4.4. First, the assembly of a UAV-derived, EMS-98-annotated validation set 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, now that the pipeline has been exercised end-to-end on a real UAV survey, a coordinated pilot deployment with a civil-protection partner is a potential next step to validate the workflow under operational conditions and to quantify the per-building classification accuracy in the UAV regime. A central scope of this research was to make the pipeline deployable for emergency actors in the field, rather than to maximise model accuracy on a benchmark. With deployability now demonstrated, a priority of future work shifts to scaling up the instance-segmentation training corpus. Taken together, these results position the MSM as a tested and 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 at scale.

CRediT authorship contribution statement

[Giovanni Galli]: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Visualization, Writing – original draft

[Marco Dubbini]: Conceptualization, Methodology, Supervision, Investigation, Writing – review & editing

[Filippo Bernardini]: Resources, Data curation, Supervision, Writing – review & editing.

[Luca Arcoraci]: Resources, Data curation, Investigation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The DFM (Database Fotografico Macrosismico) imagery underlying the training and test datasets is publicly accessible through INGV at https://dfm.ingv.it/. The pipeline source code is available at https://github.com/giolotar/Macroseismic-Survey-Mapper.

Acknowledgements

The authors gratefully acknowledge the University of Bologna (UNIBO) for institutional support of the UAV survey conducted over Piedilama (Arquata del Tronto, Central Italy). The authors also thank the Istituto Nazionale di Geofisica e Vulcanologia (INGV) for granting access to the Database Fotografico Macrosismico (DFM) used as the training corpus of this study.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author(s) used Claude (Anthropic) in order to improve the language and readability of the manuscript. After using this tool, the author(s) reviewed and edited the content as needed and take full responsibility for the content of the publication

References

[1] S. S. Matin, B. Pradhan, Challenges and limitations of earthquake-induced building damage mapping techniques using remote sensing images-A systematic review, Geocarto International 37 (21) (2022) 6186–6212. doi:10.1080/10106049.2021.1933213.
URL https://www.tandfonline.com/doi/full/10.1080/10106049.2021.1933213

[2] S. Al Shafian, D. Hu, Integrating Machine Learning and Remote Sensing in Disaster Management: A Decadal Review of Post-Disaster Building Damage Assessment, Buildings 14 (8) (2024) 2344. doi:10.3390/buildings14082344.
URL https://www.mdpi.com/2075-5309/14/8/2344

[3] G. Grünthal, European Macroseismic Scale (1998).

[4] A. Tertulliani, A. Antonucci, F. Bernardini, V. Castelli, E. Ercolani, L. Graziani, A. Maramai, M. Orlando, A. Rossi, T. Tuvè, A comprehensive integrated macroseismic dataset from multiple earthquake studies, Earth System Science Data 17 (8) (2025) 4063–4077. doi:10.5194/essd-17-4063-2025.
URL https://essd.copernicus.org/articles/17/4063/2025/

[5] A. Tertulliani, L. Arcoraci, M. Berardi, F. Bernardini, R. Camassi, C. Castellano, S. Del Mese, E. Ercolani, L. Graziani, I. Leschiutta, A. Rossi, M. Vecchi, An application of EMS98 in a medium-sized city: The case of LAquila (Central Italy) after the April 6, 2009 Mw 6.3 earthquake, Bulletin of Earthquake Engineering 9 (1) (2011) 67–80. doi:10.1007/s10518-010-9188-4.
URL http://link.springer.com/10.1007/s10518-010-9188-4

[6] D. Jozi, N. Shirzad-Ghaleroudkhani, G. Luhadia, S. Abtahi, M. Gül, Rapid post-disaster assessment of residential buildings using Unmanned Aerial Vehicles, International Journal of Disaster Risk Reduction 111 (2024) 104707. doi:10.1016/j.ijdrr.2024.104707.
URL https://www.sciencedirect.com/science/article/pii/S2212420924004692

[7] A. Rossi, A. Tertulliani, R. Azzaro, L. Graziani, A. Rovida, A. Maramai, V. Pessina, S. Hailemikael, G. Buffarini, F. Bernardini, R. Camassi, S. Del Mese, E. Ercolani, A. Fodarella, M. Locati, G. Martini, A. Paciello, S. Paolini, L. Arcoraci, C. Castellano, V. Verrubbi, M. Stucchi, The 20162017 earthquake sequence in Central Italy: macroseismic survey and damage scenario through the EMS-98 intensity assessment, Bulletin of Earthquake Engineering 17 (5) (2019) 2407–2431. doi:10.1007/s10518-019-00556-w.
URL http://link.springer.com/10.1007/s10518-019-00556-w

[8] D. Brunner, G. Lemoine, L. Bruzzone, Earthquake Damage Assessment of Buildings Using VHR Optical and SAR Imagery, IEEE Transactions on Geoscience and Remote Sensing 48 (5) (2010) 2403–2420. doi:10.1109/TGRS.2009.2038274.
URL http://ieeexplore.ieee.org/document/5411791/

[9] E. A. Ainscoe, R. Swaminathan, L. Way, S. Modugno, S. T. Chin, N. Panta, T. Crevoisier, S.-H. Yun, Earthquake damage mapped more comprehensively and accurately by radar satellites than optical imagery, Communications Earth & Environment 6 (1) (2025) 631. doi:10.1038/s43247-025-02623-4.
URL https://www.nature.com/articles/s43247-025-02623-4

[10] N. Kerle, F. Nex, M. Gerke, D. Duarte, A. Vetrivel, UAV-Based Structural Damage Mapping: A Review, ISPRS International Journal of Geo-Information 9 (1) (Dec. 2019). doi:10.3390/ijgi9010014.
URL https://www.mdpi.com/2220-9964/9/1/14

[11] A. Khan, S. Gupta, S. K. Gupta, Emerging UAV technology for disaster detection, mitigation, response, and preparedness, Journal of Field Robotics 39 (6) (2022) 905–955, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.22075. doi:10.1002/rob.22075.
URL https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.22075

[12] M. Tavakol Sadrabadi, J. Peiro, M. S. Innocente, G. Rein, Conceptual Design of a Wildfire Emergency Response System Empowered by Swarms of Unmanned Aerial Vehicles (2024). doi:10.2139/ssrn.4942867.
URL https://www.ssrn.com/abstract=4942867

[13] J. Fernandez Galarreta, N. Kerle, M. Gerke, UAV-based urban structural damage assessment using object-based image analysis and semantic reasoning, Natural Hazards and Earth System Sciences 15 (6) (2015) 1087–1101. doi:10.5194/nhess-15-1087-2015.
URL https://nhess.copernicus.org/articles/15/1087/2015/

[14] Y. Bai, H. Sezen, A. Yilmaz, End-to-end Deep Learning Methods for Automated Damage Detection in Extreme Events at Various Scales, in: 2020 25th International Conference on Pattern Recognition (ICPR), IEEE, Milan, Italy, 2021, pp. 6640–6647. doi:10.1109/ICPR48806.2021.9413041.
URL https://ieeexplore.ieee.org/document/9413041/

[15] R. Zou, J. Liu, H. Pan, D. Tang, R. Zhou, An Improved Instance Segmentation Method for Fast Assessment of Damaged Buildings Based on Post-Earthquake UAV Images, Sensors 24 (13)

(2024) 4371. doi:10.3390/s24134371.
URL https://www.mdpi.com/1424-8220/24/13/4371

[16] K. Janowicz, S. Gao, G. McKenzie, Y. Hu, B. Bhaduri, GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond, International Journal of Geographical Information Science 34 (4) (2020) 625–636. doi:10.1080/13658816.2019.1684500.

[17] J. Kallas, R. Napolitano, 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 (2025) 105358. doi:10.1016/j.ijdrr.2025.105358.
URL https://www.sciencedirect.com/science/article/pii/S2212420925001827

[18] P. Galli, S. Castenetto, E. Peronace, The MCS macroseismic survey of the Emilia 2012 earthquakes, Annals of Geophysics 55 (4) (2012) 38. doi:10.4401/ag-6163.
URL https://www.annalsofgeophysics.eu/index.php/annals/article/view/6163

[19] M. Stucchi, C. Meletti, V. Montaldo, H. Crowley, G. M. Calvi, E. Boschi, Seismic Hazard Assessment (20032009) for the Italian Building Code, Bulletin of the Seismological Society of America 101 (4) (2011) 1885–1911. doi:10.1785/0120100130.
URL https://doi.org/10.1785/0120100130

[20] A. Calantropio, F. Chiabrando, M. Codastefano, E. Bourke, DEEP LEARNING FOR AUTOMATIC BUILDING DAMAGE ASSESSMENT: APPLICATION IN POST-DISASTER SCENARIOS USING UAV DATA, ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences V-1-2021 (2021) 113–120, conference Name: XXIV ISPRS Congress <q>Imaging today, foreseeing tomorrow</q>, Commission I - 2021 edition, 5–9 July 2021. doi:10.5194/isprs-annals-V-1-2021-113-2021.
URL https://isprs-annals.copernicus.org/articles/V-1-2021/113/2021/isprs-annals-V-1-2021-113-2021.html

[21] M. Gerke, N. Kerle, Automatic Structural Seismic Damage Assessment with Airborne Oblique Pictometryl Imagery, Photogrammetric Engineering & Remote Sensing 77 (9) (2011) 885–898. doi:10.14358/PERS.77.9.885.

[22] R. Gupta, R. Hosfelt, S. Sajeev, N. Patel, B. Goodman, J. Doshi, E. Heim, H. Choset, M. Gaston, xBD: A Dataset for Assessing Building Damage from Satellite Imagery, arXiv:1911.09296 [cs] (Nov. 2019). doi:10.48550/arXiv.1911.09296.
URL http://arxiv.org/abs/1911.09296

[23] D. Duarte, F. Nex, N. Kerle, G. Vosselman, TOWARDS A MORE EFFICIENT DETECTION OF EARTHQUAKE INDUCED FAÇADE DAMAGES USING OBLIQUE UAV IMAGERY, The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-2-W6 (2017) 93–100, conference Name: International Conference on Unmanned Aerial Vehicles in Geomatics (Volume XLII-2/W6) - 4–7 September 2017, Bonn, Germany. doi:10.5194/isprs-archives-XLII-2-W6-93-2017.
URL https://isprs-archives.copernicus.org/articles/XLII-2-W6/93/2017/isprs-archives-XLII-2-W6-93-2017.html

[24] C.-S. Cheng, L. Luo, S. Murphy, Y.-C. Lee, F. Leite, A framework to enhance disaster debris estimation with AI and aerial photogrammetry, International Journal of Disaster Risk Reduction 107 (2024) 104468. doi:10.1016/j.ijdrr.2024.104468.
URL https://www.sciencedirect.com/science/article/pii/S2212420924002309

[25] J. Cheng, C. Deng, Y. Su, Z. An, Q. Wang, Methods and datasets on semantic segmentation for Unmanned Aerial Vehicle remote sensing images: A review, ISPRS Journal of Photogrammetry and Remote Sensing 211 (2024) 1–34. doi:10.1016/j.isprsjprs.2024.03.012.
URL https://linkinghub.elsevier.com/retrieve/pii/S0924271624000844

[26] J. Kim, F. Leite, A unified framework for automated damage assessment in post-disaster built environments using LiDAR point clouds, International Journal of Disaster Risk Reduction 136 (2026) 106069. doi:10.1016/j.ijdrr.2026.106069.
URL https://www.sciencedirect.com/science/article/pii/S2212420926000816

[27] G. Jocher, J. Qiu, Ultralytics YOLO11, software, version 11.0.0 (2024).
URL https://github.com/ultralytics/ultralytics

[28] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You Only Look Once: Unified, Real-Time Object Detection, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, 2016, pp. 779–788. doi:10.1109/CVPR.2016.91.

[29] L. Xing, X. Fan, Y. Dong, Z. Xiong, L. Xing, Y. Yang, H. Bai, C. Zhou, Multi-UAV cooperative system for search and rescue based on YOLOv5, International Journal of Disaster Risk Reduction 76 (2022) 102972. doi:10.1016/j.ijdrr.2022.102972.
URL https://www.sciencedirect.com/science/article/pii/S2212420922001911

[30] B. Kim, S. Cho, Automated Multiple Concrete Damage Detection Using Instance Segmentation Deep Learning Model, Applied Sciences 10 (22) (2020) 8008. doi:10.3390/app10228008.
URL https://www.mdpi.com/2076-3417/10/22/8008

[31] I. Bouchard, M.-. Rancourt, D. Aloise, F. Kalaitzis, On Transfer Learning for Building Damage Assessment from Satellite Imagery in Emergency Contexts, Remote Sensing 14 (11) (May 2022). doi:10.3390/rs14112532.
URL https://www.mdpi.com/2072-4292/14/11/2532

[32] S. Wiguna, B. Adriano, E. Mas, S. Koshimura, 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 (2024) 5651–5667. doi:10.1109/JSTARS.2024.3367853.
URL https://ieeexplore.ieee.org/document/10440409/

[33] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, Q. He, A Comprehensive Survey on Transfer Learning, arXiv:1911.02685 [cs] (Jun. 2020). doi:10.48550/arXiv.1911.02685.
URL http://arxiv.org/abs/1911.02685

[34] W. Yang, X. Zhang, P. Luo, Transferability of Convolutional Neural Network Models for Identifying Damaged Buildings Due to Earthquake, Remote Sensing 13 (3) (Jan. 2021). doi:10.3390/rs13030504.
URL https://www.mdpi.com/2072-4292/13/3/504

[35] M. Kakooei, Y. Baleghi, Fusion of satellite, aircraft, and UAV data for automatic disaster damage assessment, International Journal of Remote Sensing 38 (8-10) (2017) 2511–2534, _eprint: https://doi.org/10.1080/01431161.2017.1294780. doi:10.1080/01431161.2017.1294780.
URL https://doi.org/10.1080/01431161.2017.1294780

[36] D. Song, X. Tan, B. Wang, L. Zhang, X. Shan, J. Cui, 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) (2020) 1040–1066, _eprint: https://doi.org/10.1080/01431161.2019.1655175. doi:10.1080/01431161.2019.1655175.
URL https://doi.org/10.1080/01431161.2019.1655175

[37] C. Meletti, V. Montaldo, M. Stucchi, F. Martinelli, Database della pericolosità sismica MPS04 (2006). doi:10.13127/SH/MPS04/DB.
URL http://esse1.mi.ingv.it/