Real Time Earthquake Detection Using YOLOv8 and Spectrogram Analysis

Mahmoud Salam

National Research Institute of Astronomy and Geophysics (NRIAG)

INTRODUCTION AND MAIN RESULTS

Traditional earthquake detection methods can be slow and prone to human error, especially in real-time monitoring scenarios. Rapid and accurate detection is crucial for early warning systems. This research proposes a novel AI-powered approach using the state-of-the-art YOLOv8 object detection model to automatically and instantly identify earthquakes in seismic data. Leveraging computer vision for seismic analysis allows for sub-second inference times, making it ideal for real-time deployment.

SnT 2025
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE
8 SEPTEMBER ONLINE DAY
9 TO 12 SEPTEMBER AT HOFBURG PALACE, VIENNA & ONLINE

Real Time Earthquake Detection Using YOLOv8 and Spectrogram Analysis

Mahmoud Salam

P3.5-442

Dataset & Preparation

The model was trained on the INSTANCE dataset (Michelini et al., 2021). This dataset comprises approximately 1.2 million three-component waveform traces from about 50,000 earthquakes and over 130,000 noise recordings, from the Italian region between 2005 and 2020. Sourced from the Italian National Institute of Geophysics and Volcanology (INGV), the dataset includes events with magnitudes from 0.0 to 6.5. All waveforms are standardized to 120-second lengths at a 100 Hz sampling rate .

The process begins with the transformation of raw, continuous seismic waveform data into time-frequency representations known as spectrograms. This conversion is critical, as it translates the one-dimensional seismic signal into a two-dimensional image where the distinct energy patterns of seismic events become visually discernible features against the background noise. These patterns often differ between the higher-frequency content of local events and the lower-frequency, longer-duration signals characteristic of regional earthquakes, providing a basis for subsequent classification.

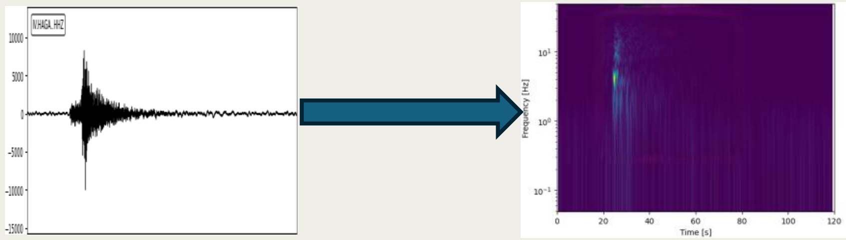

Data annotation

Following the creation of these spectrograms, **an automated annotation process** is executed to generate the labeled dataset for machine learning. This algorithm employs a series of signal processing techniques to programmatically define bounding boxes without human intervention. The left boundary of each box is the precise **reported P-phase arrival time**. The right boundary is determined by a separate algorithm that identifies the **trigger-off point**, where the signal's amplitude stabilizes back to the baseline noise level. Vertically, the lower and upper frequency boundaries of the box are defined by analyzing the Signal-to-Noise Ratio (SNR) across the frequency spectrum. The box is automatically drawn to encapsulate only the frequency band where the SNR exceeds a predefined threshold, capturing the event's specific energy signature. Finally, a classification for each bounded event as "local" or "regional" based on the reported catalog and the calculated distance. This auto-generated dataset, comprising thousands of accurately labeled images, serves as the foundational ground truth required to train a sophisticated object detection model.

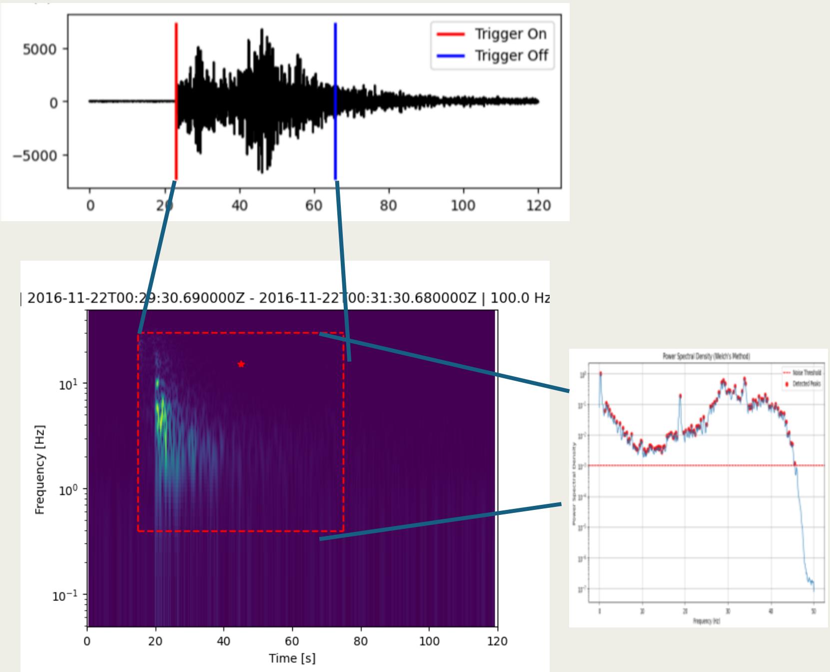

SnT 2025
8 SEPTEMBER ONLINE DAY
9 TO 12 SEPTEMBER AT HOFBURG PALACE, VIENNA & ONLINE
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE

Real Time Earthquake Detection Using YOLOv8 and Spectrogram Analysis

Mahmoud Salam

P3.5-442

Model Training

YOLOv8 is a state-of-the-art deep learning architecture renowned for its speed and accuracy in identifying and locating objects within images. During training, the model learns to correlate the visual patterns in the spectrograms—such as the shape, size, and frequency content of the signal. It learns to not only detect the presence of an event but also to classify it simultaneously, distinguishing between a local and a regional earthquake based on the learned features.

Based on the training results, the YOLOv8 model demonstrates successful learning and strong overall performance, achieving a high mean average precision (mAP50) of approximately 80%, which indicates excellent detection accuracy. The loss curves show ideal behavior, with training and validation losses decreasing steadily and converging without signs of overfitting, confirming the model generalizes well to unseen data. However, a key limitation is the moderate recall rate of around 60%, suggesting the model misses a significant portion of true events, particularly confusing regional and local classes or failing to detect weaker signals. Additionally, while precision remains high at roughly 80%, reducing false negatives and refining localization accuracy—especially for bounding box boundaries—should be prioritized to enhance reliability for real-time seismic monitoring applications.

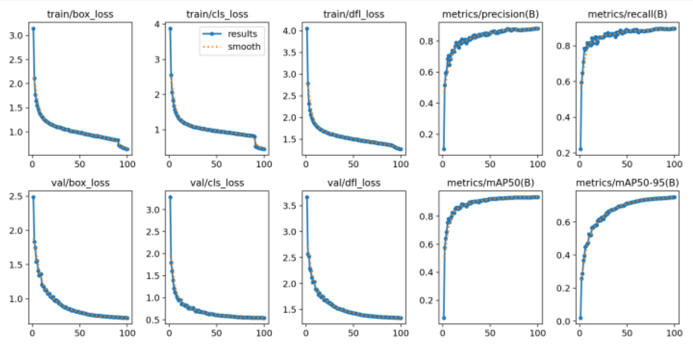

Results

The confusion matrix shows that the model performs very well in distinguishing between local and regional earthquakes, with 96% and 91% accuracy, respectively. There is minimal confusion between these two classes, which indicates the model has effectively learned to differentiate seismic patterns associated with each. However, a significant issue arises with the background class. While a concerning 71% are misclassified as regional earthquakes, and 29% of local earthquakes are misclassified as background. This misclassification could either be the result of undetected events (i.e., the model missed true earthquakes) or incorrectly labeled background (false positives where the model identifies background noise as an earthquake). To resolve this, it would be important to ensure that the model's detection capabilities are robust enough to distinguish between true seismic events and background noise, possibly through better data labeling and more training on varied background and earthquake data.

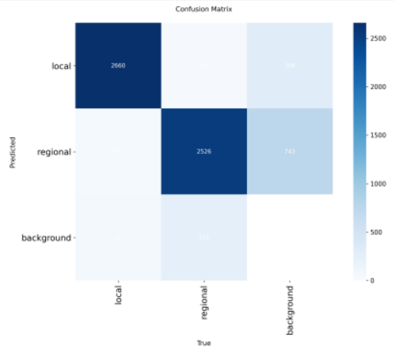

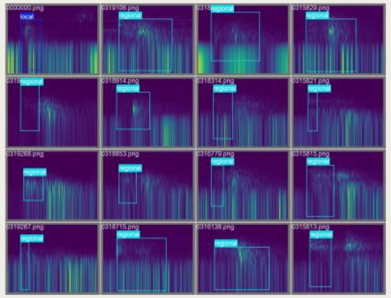

Model deployment

Finally, the trained and validated YOLOv8 model is deployed into an operational pipeline for rapid, automated analysis. In this production environment, continuous seismic data is streamed, automatically converted into spectrograms in near-real-time, and fed into the model for inference. The model scans each new spectrogram, rapidly identifies the onset and termination of seismic events, draws bounding boxes around them, and assigns a classification label—all within milliseconds. This system provides seismologists with immediate, automated detection and preliminary classification of events, significantly accelerating response times for seismic monitoring and early warning systems.