

GaLeNet: Multimodal Learning for Disaster Prediction, Management and Relief

Anonymous Author(s)

ABSTRACT

After a natural disaster, such as a hurricane, millions are left in need of emergency assistance. To allocate resources optimally, human planners need to accurately analyze data that can flow in large volumes from several sources. This motivates the development of multimodal machine learning frameworks that can integrate multiple data sources and leverage them efficiently. To date, the research community has mainly focused on unimodal reasoning to provide granular assessments of the damage. Moreover, previous studies mostly rely on post-disaster images, which may take several days to become available. In this work, we propose a multimodal framework (GaLeNet) for assessing the severity of damage by complementing pre-disaster images with weather data and the trajectory of the hurricane. Through extensive experiments on data from two hurricanes, we demonstrate (i) the merits of multimodal approaches compared to unimodal methods, and (ii) the effectiveness of GaLeNet at fusing various modalities. Furthermore, we show that GaLeNet can leverage pre-disaster images in the absence of post-disaster images, preventing substantial delays in decision making.

CCS CONCEPTS

• **Computing methodologies** → *Machine learning; Neural network*; • **Applied computing** → *Environmental sciences*.

KEYWORDS

disaster management, hurricane, neural networks, multimodal reasoning, CLIP embeddings

ACM Reference Format:

Anonymous Author(s). 2022. GaLeNet: Multimodal Learning for Disaster Prediction, Management and Relief. In *Proceedings of KDD'22 Fragile Earth Workshop*. ACM, New York, NY, USA, 8 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

In Haiti alone, Hurricane Matthew left an estimated 180,000 people homeless and 1.4 million in need of emergency assistance [25]. When so many people are affected, it is critical to get an accurate assessment of the location and severity of damage so that resources can be allocated quickly to where they are needed most. The earlier this picture of destruction can be assembled, the better a natural

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD'22 Fragile Earth Workshop, August 14–18, 2022, Washington DC, USA

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

disaster can be managed. Disaster management and prediction could become even more important in the future, as the frequency of natural disasters might be increasing due to climate change [7].

One way to identify affected areas is to use post-disaster (“gray skies”) images captured by satellite, drone, or plane in the immediate aftermath of a disaster. However, there can be a delay of three days before such images are available [29]. There is a multitude of data that can be gathered before or during the event such as pre-disaster (“blue skies”) images, weather features or dynamics, and the trajectory of the hurricane. Are we able to reason over such a diverse range of data modalities to predict which buildings will be affected most before we have direct visual evidence of the destruction?

Recently, several works have been published on applying machine learning to natural disaster prediction and management by predicting the location and severity of damage to buildings [5, 13, 21, 26, 31–33]. However, these rely on direct visual evidence of the damage provided by the post-disaster images. There are also an increasing number of publications on multimodal machine learning, but these use modalities that are naturally aligned (e.g. twitter images and text) and typically only provide “big picture” [3] insights about the damage such as classifying image-text pairs from social media for their informativeness, type of emergency and severity of emergency [1].

In this work, we propose a multimodal framework to predict the severity of the damage *before* images of the damage are available. The motivation is to allow for early triage and allocation of resources. We also show how the accuracy of the framework improves once the post-disaster images are collected. This could allow authorities to adapt their relief program as new information becomes available.

Contributions. To our knowledge, this work is the first to attempt using multimodal machine learning to build a granular picture of damage severity without using post-disaster imaging. This is achieved by: (i) aligning the data modalities via the longitude, latitude and time of event, and (ii) supplementing pre-disaster images with weather data as well as information about the hurricane trajectory through a novel featurization method. The framework also allows for the assessment to be updated once post-disaster imagery is available. Through the use of pre-trained image embeddings, we are able to efficiently train our framework on a limited set of examples. We call our framework GaLeNet¹.

2 RELATED WORK

The need for rapid damage assessment following a natural disaster has driven research into machine learning methods that can quickly process data to assess the location and severity of damage [3]. Decision makers are often faced with the challenge of promptly

¹Coined by conjoining the words “Gale” and “LeNet”.

integrating several streams of information, reasoning over them, and coming to an accurate assessment [10] which may evolve as more data becomes available.

Multimodal learning is a natural solution to this problem; however, work in this area is still limited to understanding damage at a high level [3]. Methods capable of making granular inferences have so far been limited to unimodal image-based methods that are mostly single task [5, 21, 26, 31–33] and sometimes multitask [13] - all of which rely on the post-disaster image.

Despite the recent advances in deep learning, the applications of these models for natural disaster prediction is still limited due to the scarcity of training data. This is even more severe in a multimodal setting where several data sources have to be sourced and merged. This has limited the application of deep learning, although these methods are generally effective given enough data [4]. Here, we overcome this problem by using pre-trained models as feature extractors (for images) and handcrafted representations of the hurricane trajectory and weather data leading up to the event. For images, we observe that using multiple scales helps with the representation, an observation which motivated [5].

For the hurricane trajectory, we present a novel featurization method based on calculating the closest point along the hurricane trajectory. The authors are also not aware of any building damage prediction models that use the hurricane trajectory as an input modality.

Aligning separate streams of data is one of the challenges of multimodal learning [6]. In this work, we align data using the geolocation (longitude and latitude) and timestamp of the disaster.

3 METHODS

3.1 Problem Statement & Scope

We consider both a *proactive* and a *reactive* scenario. In the *reactive* case, the goal is to predict which buildings will be damaged, and the severity of the damage, using direct visual evidence collected through post-disaster satellite imaging. However, in the *proactive* case, we are only allowed to use contextual information such as pre-disaster imaging or weather data.

In this work, we limit the scope to include buildings in the vicinity of hurricanes Matthew (2016) [27] and Michael (2018) [8]. We focus on hurricanes as there is more contextual information (e.g. hurricane trajectory or weather data) than is available for other disasters such as forest fires, thus making the problem more compatible with multimodal learning.

We treat the problem as a classification task, where for each building, we classify the damage into one of four levels of severity: (i) no damage, (ii) minor damage, (iii) major damage, and (iv) destroyed. The labels and class definitions are taken from the xBD dataset [12]. Other than selecting only examples related to wind damage due to hurricanes, we used the original train, test, and hold splits from the xView 2 challenge, throughout.

3.2 Data Modalities

The xBD dataset [12] was used as the foundation of our multimodal dataset, as it contains labels for the severity of damage as well as the geolocation and time of the disasters. We used these as anchor points to align all of our different modalities: (i) pre-disaster satellite

images, (ii) post-disaster satellite images, (iii) weather data, and (iv) the hurricane trajectory.

After filtering the original xBD dataset [12] for hurricanes and wind damage, 36,625, 9,283, and 12,791 building were identified for training, validation and testing. The dataset is somewhat imbalanced with 48% of the labels being no damage and the remaining 52% consisting of the other three levels of damage (33% minor damage, 11% major damage, and 8% destroyed).

3.2.1 Image Data.

Source. The images from the xBD dataset [12] were pre-processed by taking the centroid of each building polygon and then taking a centered crop, as opposed to using the full satellite image.

Representation. We tried both uniscale and multiscale approaches for visual representation. A zoomed out view (denoted by Scale-1x) was created by scaling the image by $s \approx 11/\sqrt{A}$ (where A is the area of the building polygon) in each dimension before performing a crop, centered around the building center, of 224×224 pixels. Additional crops were performed by increasing the scale to $4s$ (Scale-4x), $16s$ (Scale-16x), and $32s$ (Scale-32x) - progressively zooming in on the building.

Visual features were extracted using: (i) a pre-trained CNN baseline, (ii) the CLIP [24] visual branch, and (iii) a pre-trained U-Net on satellite images to segment buildings by type (houses, building, and sheds/garages) [20]. The CNN baseline consists of five convolutional blocks followed by two fully-connected layers. We trained this model on the cropped images from the xBD training set for the damage classification task. We explored four different versions of CLIP [24] for extracting visual features: (i) CLIP ViT-B/32 (ii) CLIP ViT-B/16 (iii) CLIP ViT-L/14 and (iv) CLIP ViT-L/14@336px. For the U-Net, we used the pre-trained weights available from [20].

3.2.2 Hurricane Trajectory.

Source. The trajectory of the two hurricanes were extracted from reports published by the NOAA [8, 27]. The trajectory is a series of longitude and latitude coordinates with an associated time, wind speed and central pressure.

Representation. To featurize the hurricane trajectory, we used the haversine formula to compute the shortest distance between each building and the trajectory. This shortest distance (in km) and the wind speed and pressure at this point in the trajectory were extracted. In case of ties (i.e. multiple points on the trajectory that had the same distance to the building), the maximum wind speed and pressure were used. We also experimented with more complex multiscale and multipoint schemes as well as distance-weighted features, but these offered no benefit.

3.2.3 Weather Data.

Source. The weather data were extracted from OpenWeather² and consists of many factors including, but not limited to, the temperature, wind speed and direction, precipitation, humidity, pressure, and visibility. We collected daily weather data for seven days preceding each hurricane, with eight time points per day.

²Data was collected using the OpenWeather API <https://openweathermap.org/>

Representation. The best representation of the weather data was simply the average of each weather feature across all time points. We also tried using off-the-shelf feature engineering libraries for time series, such as TSFresh [11] and Catch22 [19], but these offered limited benefit and resulted in a substantial increase in the number of input features and model parameters, making it prone to overfitting.

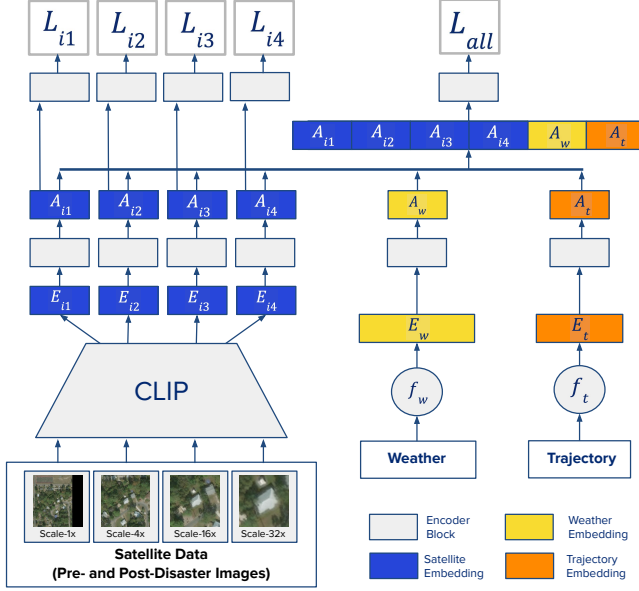


Figure 1: Schematic diagram of our proposed framework

3.3 Metrics

The following metrics were used for evaluation: (i) Area Under the Precision Recall Curve (PR AUC), (ii) Area Under the ROC Curve (ROC AUC), and (iii) Balanced Accuracy (Bal. Acc.) which is the unweighted average of recall obtained on each class. Since we are dealing with an imbalanced dataset, we use macro averaging to obtain the PR AUC and ROC AUC curves.

4 EXPERIMENTS

To study the contribution of each modality, we conduct two sets of experiments. First, we evaluate and compare the representations of each modality separately on damage classification. Then, using the best representation of each modality, we assess the effectiveness of our multimodal framework for both the proactive and reactive scenarios.

4.1 Single Modalities

To evaluate and compare the different modalities and their representations, we follow the common practice [16] and train a linear logistic regression (LogReg) model for damage severity classification. We train LogReg using the L-BFGS [18] solver. We choose the best value for the inverse regularization strength C by running a grid search over $C \in 10^{[-3,3]}$ on the validation data.

4.2 Multiple Modalities

We assess the effectiveness of our multimodal framework for proactive and reactive scenarios by comparing it against two models: (i) LogReg, and (ii) Concat-MLP. For LogReg, we use the setup described in 4.1. In the reactive case, we use the hurricane trajectory, weather data, and the post-disaster satellite images. In the proactive case, however, we use pre-disaster images instead of the post-disaster images.

Concat-MLP. We begin by concatenating the best representations of each modality, namely, the embeddings extracted from CLIP ViT-L/14@336px corresponding to the four different image scales, denoted by $\{E_{i1}, E_{i2}, E_{i3}, E_{i4}\}$, the weather embedding E_w , and the trajectory embedding E_t . Let $E_{all} = [E_{i1}, E_{i2}, E_{i3}, E_{i4}, E_w, E_t]$ be the concatenated representation. Next, E_{all} is passed through a network of two fully-connected layers. The first layer contains 128 nodes, followed by 32 nodes in the second layer, and both layers use ReLU [2] activation. The output of the network is fed into a softmax layer for classification.

GaLeNet. We observe that an early-stage naïve concatenation of feature embeddings, extracted from different modalities, can lead to difficulties in training due to: (i) different modalities having variable rates of learning [30], and (ii) modalities containing different levels of information for the task at hand. To combat these problems, we use *late-fusion*; a common paradigm in which modality-specific latent representations are learned to aid the process of fusion [14, 23, 34].

To this end, our proposed framework (Figure 1) jointly trains multiple modality-specific encoders, whose intermediate activations are then concatenated. Each encoder follows a similar sequence of computations; *linear projection* \rightarrow *batch normalization* \rightarrow *ReLU* \rightarrow *dropout*. The CLIP embeddings $\{E_{i1}, E_{i2}, E_{i3}, E_{i4}\} \in \mathbb{R}^{56}$. Similarly, the encoders corresponding to E_w and E_t produce $A_w \in \mathbb{R}^{16}$ and $A_t \in \mathbb{R}^3$ respectively. Next, these intermediate activations are concatenated, yielding $A_{all} \in \mathbb{R}^{243}$, and passed through a fusion encoder before finally being fed into a classification layer. Motivated by [28], each of $\{A_{i1}, A_{i2}, A_{i3}, A_{i4}\}$ is additionally fed into its respective classification layer to accelerate the optimization process.

Combining all the losses computed from the classification layers, the overall optimization objective becomes

$$L = \arg \min_{\theta} \sum_{j=1}^4 L_{ij} + L_{all} \quad (1)$$

where L_{ij} corresponds to the losses computed from $\{A_{i1}, A_{i2}, A_{i3}, A_{i4}\}$, L_{all} corresponds to the loss computed from A_{all} , and θ is the set of learnable parameters.

Training Configuration. To ensure consistency between Concat-MLP and GaLeNet, we employ Focal Loss [17] due to its superior performance on imbalanced datasets. We use Adam optimizer [15], initialized with a learning rate of $1e-4$. To combat overfitting, we use Early Stopping [9] with the patience set to 5. Finally, we run each model 5 times with random initialization of weights and report the averaged metrics.

5 RESULTS

5.1 Single Modalities

It took some optimization to find the best representation of the image data. Table 1 shows the performance of the different feature extraction approaches at a fixed scale (Scale-32x). In general, the CLIP representation outperformed both the CNN and the U-Net approaches. This is interesting given that CLIP has not been explicitly trained to detect buildings or damage from satellite images. But, it does have the benefit of being exposed to a much larger variety and number of images during pre-training. We also observe that the performance increases from older to newer versions of CLIP (from top to bottom), which is consistent with the original generalization trend of CLIP models shown by [22]. The largest CLIP model (ViT-L/14@336px) performed best and was used for all other experiments that we report.

Table 1: Comparison of various feature extractors

Representation	Bal. Acc.	PR AUC	ROC AUC
CNN	0.5091	0.5711	0.8310
U-Net	0.3574	0.3970	0.6976
CLIP ViT-B/32	0.5352	0.5925	0.8300
CLIP ViT-B/16	0.5499	0.6023	0.8296
CLIP ViT-L/14	0.5564	0.6107	0.8404
CLIP ViT-L/14@336px	0.5684	0.6183	0.8443

Table 2 presents the performance metrics of CLIP ViT-L/14@336px embeddings obtained using various cropping strategies on pre- and post-disaster satellite images. Interestingly, we notice a difference between the pre- and post-disaster representations. In the pre-disaster case, the performance improves progressively as we “zoom out” from the building. However, in the post-disaster case the opposite is true and “zooming in” to the building seems to boost performance. We reason that this is because in the post-disaster case the direct visual evidence of the building damage is important to assessing damage severity. However, in the pre-disaster case this is not available, so the model relies on inferring potential damage using the contextual information surrounding the building. This might include whether it is sheltered by other buildings or forests, or whether it is close to loose structures that could get uplifted and blown into the building.

In both cases, the highest performance is achieved by concatenating all the different scales, implying that both pre- and post-disaster damage assessment benefit from combining information obtained at multiple scales.

Table 3 (rows 1-4) show the performance of each modality in the unimodal setting using LogReg as the model. The ROC AUC performance shows that on their own, each modality provides some utility for assessing building damage severity, even for proactive modalities that do have access to direct visual evidence of the damage. As expected, the post-disaster image modality provides a significant boost in performance over the proactive modalities.

Figure 2(a-d) shows ROC curves for the unimodal baselines. It is evident that for the proactive modalities (a-c), there is a noticeable drop in performance for the “Major Damage” class, however this is

not the case for the post-disaster image modality. When studying the class predictions of the model, it is evident that the proactive models tend to confuse “Major Damage” with “Destroyed”. In the post-disaster case, the model becomes much better at distinguishing the different levels of damage severity as it has direct visual evidence of the damage.

5.2 Multiple Modalities

Table 3 (rows 5-10) provides a comparison of our proposed framework (GaLeNet) to the LogReg and Concat-MLP baselines for both the proactive and reactive scenarios.

Firstly, we observe that the LogReg is able to benefit from using all input modalities, as the multimodal LogReg outperforms all of the unimodal LogReg models in both the proactive and reactive cases. This implies that the different modalities have some complementary information for inferring the damage severity, even if provided with post-disaster images.

Secondly, we note that a naïve neural network approach (Concat-MLP) does not offer much benefit over the LogReg, as it does not consistently outperform the LogReg in either the proactive or reactive cases.

GaLeNet outperforms both fusion baselines across all metrics for both scenarios. Focusing on the Proactive use-case, GaLeNet achieves an increase of ~14% in Bal. Acc. over LogReg. In the reactive use case, GaLeNet achieves similar performance boost in Bal. Acc.

Qualitatively, GaLeNet demonstrates its capability to correctly identify the severity of damage across buildings of multiple sizes (Figures 3 and 4). This can be attributed to the availability of visual features extracted at different image scales and the normalization of the image scale.

Finally, it is worth noting that GaLeNet outperforms its closest architectural neighbor, Concat-MLP, across all metrics while having only 189K parameters (52.75% fewer parameters compared to Concat-MLP). This further attests to the effectiveness of GaLeNet’s compact architectural design.

6 LIMITATIONS & FUTURE WORK

During the process of this work, we found several limitations which open up impactful avenues of future work:

- We experimented with additional modalities such as elevation data and twitter data, and while these gave above chance performance in an unimodal setting, they did not offer any benefit in a multimodal setting. This could be due to the limited size of our dataset, we had to use a relatively shallow model which was unable to reason across these additional modalities. In addition, the twitter dataset was relatively sparse, and the resultant model suffered from “modality dropout”. We feel that a solution to this is essential to real-world application of such a model.
- Several attempts were made to improve the representation of the weather data and the hurricane trajectory data, however, it was found that only relatively simple featurizations of each offered the best performance. We attribute this to the granularity (in terms of longitude and latitude) of the data and to the overall size of the dataset and the inability

Table 2: Comparison of various cropping strategies and performance for pre- and post-disaster images.

Image Scale	Pre-disaster			Post-disaster		
	Bal. Acc.	PR AUC	ROC AUC	Bal. Acc.	PR AUC	ROC AUC
Scale-1x	0.4963	0.5329	0.7878	0.4929	0.5598	0.8129
Scale-4x	0.4835	0.5191	0.7797	0.5147	0.5804	0.8229
Scale-16x	0.4762	0.5201	0.7799	0.5381	0.6083	0.8386
Scale-32x	0.4588	0.4944	0.7620	0.5684	0.6183	0.8443
All Scales	0.4931	0.5439	0.7962	0.5707	0.6430	0.8570

Table 3: Comparison of unimodal and multimodal baselines with GaLeNet, for the modalities: weather data (W), hurricane trajectory (T), pre-disaster image (Pre), and post-disaster image (Post).

Model	Features	Scenario	Bal. Acc.	PR AUC	ROC AUC
LogReg	W	Proactive	0.5238	0.4369	0.7122
LogReg	T	Proactive	0.4348	0.5145	0.7508
LogReg	Pre	Proactive	0.4931	0.5439	0.7962
LogReg	Post	Reactive	0.5707	0.6430	0.8570
LogReg	W + T + Pre	Proactive	0.5110	0.5533	0.8090
Concat-MLP	W + T + Pre	Proactive	0.6357	0.5518	0.8072
GaLeNet	W + T + Pre	Proactive	0.6495	0.5645	0.8140
LogReg	W + T + Post	Reactive	0.5773	0.6472	0.8631
Concat-MLP	W + T + Post	Reactive	0.6798	0.6556	0.8648
GaLeNet	W + T + Post	Reactive	0.6875	0.6680	0.8732

of the model to learn a more nuanced representation of the data. The dataset size also prevented us experimenting with different co-learning strategies.

- Following error analysis, we realized that our framework struggles to identify damage on buildings that are circular. This can be attributed to their rare occurrence in the real world, and consequently in the dataset. Future work can, therefore, explore extending existing datasets to cover edge cases.
- The dataset only contained two hurricanes, so we were unable to test generalization to new natural disasters. Collecting data and testing across more natural disasters is an essential next step (in our roadmap) before these models can be utilized in the real world.

7 CONCLUSION

We proposed a multimodal framework, GaLeNet, that is capable of predicting the severity of damage to buildings after a hurricane even if it does not have access to direct visual evidence of the damage. Moreover, when such information is available, GaLeNet is able to use it to increase the accuracy of its predictions. Through extensive evaluation with data from two hurricanes, we show the effectiveness of GaLeNet by comparing it against multiple unimodal and multimodal baselines.

We believe that such a framework could provide crucial early insights and intelligence on the natural disaster before the true damage is known. Thus, it could help allocate resources for disaster

relief where they matter most, and update this allocation as more data becomes available.

In addition, it is possible that proactive data such as weather predictions, predicted hurricane trajectories and “blue skies” imaging could provide useful insights about the risk of particular buildings to damage in the event of a natural disaster. Such actionable insights could be used to strengthen our defenses against disasters before they occur.

As climate change and its effects become more pronounced, the need to address its challenges has never been more pressing. We hope that GaLeNet, despite its limitations, can ignite future research and development in this direction.

REFERENCES

- [1] Mahdi Abavisani, Liwei Wu, Shengli Hu, Joel Tetreault, and Alejandro Jaimes. 2020. Multimodal categorization of crisis events in social media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14679–14689.
- [2] Abien Fred Agarap. 2018. Deep Learning using Rectified Linear Units (ReLU). *CoRR* abs/1803.08375 (2018). arXiv:1803.08375 <http://arxiv.org/abs/1803.08375>
- [3] Nilani Algiriyage, Raj Prasanna, Kristin Stock, Emma EH Doyle, and David Johnston. 2022. Multi-source Multimodal Data and Deep Learning for Disaster Response: A Systematic Review. *SN Computer Science* 3, 1 (2022), 1–29.
- [4] Rania Rizki Arinta and Emanuel Andi WR. 2019. Natural disaster application on big data and machine learning: A review. In *2019 4th International Conference on Information Technology, Information Systems and Electrical Engineering (ICITISEE)*. IEEE, 249–254.
- [5] Yanbing Bai, Junjie Hu, Jinhua Su, Xing Liu, Haoyu Liu, Xianwen He, Shengwang Meng, Erick Mas, and Shunichi Koshimura. 2020. Pyramid pooling module-based semi-siamese network: A benchmark model for assessing building damage from xBD satellite imagery datasets. *Remote Sensing* 12, 24 (2020), 4055.

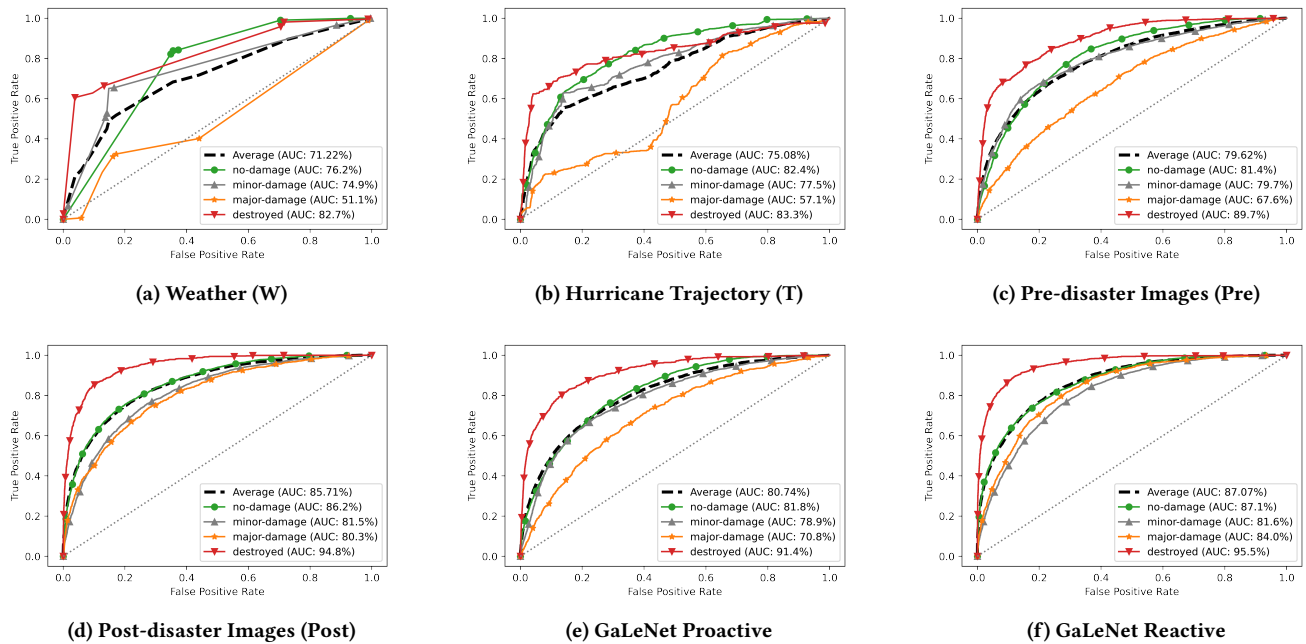



Figure 2: Comparison of ROC curves between unimodal baselines (a-d) and GaLeNet used for the proactive (e) and reactive (f) use cases.


- [6] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* 41, 2 (2018), 423–443.
- [7] Sandra Banholzer, James Kossin, and Simon Donner. 2014. The impact of climate change on natural disasters. In *Reducing disaster: Early warning systems for climate change*. Springer, 21–49.
- [8] John L Beven II, Robbie Berg, and Andrew Hagen. 2018. Hurricane Michael. Retrieved May 23, 2022 from https://www.nhc.noaa.gov/data/tcr/AL142018_Michael.pdf
- [9] Rich Caruana, Steve Lawrence, and Lee Giles. 2000. Overfitting in Neural Nets: Backpropagation, Conjugate Gradient, and Early Stopping (*NIPS'00*). MIT Press, Cambridge, MA, USA, 381–387.
- [10] Vinay Chamola, Vikas Hassija, Sakshi Gupta, Adit Goyal, Mohsen Guizani, and Biplab Sikdar. 2020. Disaster and pandemic management using machine learning: a survey. *IEEE Internet of Things Journal* 8, 21 (2020), 16047–16071.
- [11] Maximilian Christ, Nils Braun, Julius Neuffer, and Andreas W Kempa-Liehr. 2018. Time series feature extraction on basis of scalable hypothesis tests (tsfresh—a python package). *Neurocomputing* 307 (2018), 72–77.
- [12] Ritwik Gupta, Richard Hosfelt, et al. 2019. xbd: A dataset for assessing building damage from satellite imagery. *arXiv preprint arXiv:1911.09296* (2019).
- [13] Rohit Gupta and Mubarak Shah. 2021. Rescuenet: Joint building segmentation and damage assessment from satellite imagery. In *ICPR*.
- [14] Cheng Ju, Aurelien Bibaut, and Mark Laan. 2017. The Relative Performance of Ensemble Methods with Deep Convolutional Neural Networks for Image Classification. *Journal of Applied Statistics* 45 (04 2017). <https://doi.org/10.1080/02664763.2018.1441383>
- [15] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- [16] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. 2019. Revisiting self-supervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1920–1929.
- [17] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. *CoRR* (2017).
- [18] Dong C Liu and Jorge Nocedal. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical programming* 45, 1 (1989), 503–528.
- [19] Carl H Lubba, Sarab S Sethi, Philip Knaute, Simon R Schultz, Ben D Fulcher, and Nick S Jones. 2019. catch22: Canonical time-series characteristics. *Data Mining and Knowledge Discovery* 33, 6 (2019), 1821–1852.
- [20] obravo7. 2021. obravo7/satellite-segmentation-pytorch: Multi-class satellite semantic segmentation using PYTORCH framework. Retrieved May 23, 2022 from <https://github.com/obravo7/satellite-segmentation-pytorch>
- [21] Victor Oludare, Landry Kezebou, Karen Panetta, and Sos Agaian. 2021. Semi-supervised learning for improved post-disaster damage assessment from satellite imagery. In *Multimodal Image Exploitation and Learning 2021*, Vol. 11734. International Society for Optics and Photonics, 117340O.
- [22] OpenAI. 2021. Clip: Connecting text and images. Retrieved May 23, 2022 from <https://openai.com/blog/clip/>
- [23] Andrew Owens and Alexei A. Efros. 2018. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. *CoRR* abs/1804.03641 (2018). <http://arxiv.org/abs/1804.03641>
- [24] Alec Radford et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [25] George Rodriguez. 2016. Aftermath of Hurricane Matthew Brings Threat of Food Shortages, More Cholera. (2016).
- [26] Yu Shen, Sijie Zhu, Taojiannan Yang, and Chen Chen. 2020. Cross-directional feature fusion network for building damage assessment from satellite imagery. *arXiv preprint arXiv:2010.14014* (2020).
- [27] Stewart R Stacy. 2016. Hurricane Matthew. Retrieved May 23, 2022 from https://www.nhc.noaa.gov/data/tcr/AL142016_Matthew.pdf
- [28] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [29] Stefan Voigt, Thomas Kemper, et al. 2007. Satellite Image Analysis for Disaster and Crisis-Management Support. *IEEE T. Geoscience and Remote Sensing* 45 (06 2007), 1520–1528. <https://doi.org/10.1109/TGRS.2007.895830>
- [30] Weiyao Wang, Du Tran, and Matt Feiszli. 2019. What Makes Training Multimodal Networks Hard? *CoRR* (2019). <http://arxiv.org/abs/1905.12681>
- [31] Ethan Weber and Hassan Kané. 2020. Building disaster damage assessment in satellite imagery with multi-temporal fusion. *arXiv preprint arXiv:2004.05525* (2020).
- [32] Chuyi Wu, Feng Zhang, et al. 2021. Building damage detection using U-Net with attention mechanism from pre-and post-disaster remote sensing datasets. *Remote Sensing* 13, 5 (2021), 905.
- [33] Joseph Z Xu, Wenhan Lu, Zebo Li, Pranav Khaitan, and Valeriya Zaytseva. 2019. Building damage detection in satellite imagery using convolutional neural networks. *arXiv preprint arXiv:1910.06444* (2019).
- [34] Yunke Zhang, Lixue Gong, Lubin Fan, Peiran Ren, Qixing Huang, Hujun Bao, and Weiwei Xu. 2019. A late fusion cnn for digital matting. In *CVPR*. 7469–7478.

697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754


755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812



(a) Ground truth - Hurricane Matthew



(b) GaLeNet proactive - Hurricane Matthew



(c) Ground truth - Hurricane Michael



(d) GaLeNet proactive - Hurricane Michael

Figure 3: Visual comparison between ground truth labels and GaLeNet predictions for the *proactive* scenario. Images from hurricanes Matthew and Michael are shown in (a, b) and (c, d), respectively. “No Damage“, “Minor Damage”, “Major Damage” and “Destroyed” are shown in green, wheat, orange and red, respectively. Correct predictions are filled in the relevant colour, whereas misclassified predictions are outlined in the colour of the predicted label. Note that the building masks are taken from the xBD dataset and not predicted by the model, they are used only for visualization purposes.




(a) Ground truth - Hurricane Matthew



(b) GaLeNet reactive - Hurricane Matthew



(c) Ground truth - Hurricane Michael



(d) GaLeNet reactive - Hurricane Michael

Figure 4: Visual comparison between ground truth labels and GaLeNet predictions for the *reactive* scenario. Images from hurricanes Matthew and Michael are shown in (a, b) and (c, d), respectively. “No Damage”, “Minor Damage”, “Major Damage” and “Destroyed” are shown in green, wheat, orange and red, respectively. Correct predictions are filled in the relevant colour, whereas misclassified predictions are outlined in the colour of the predicted label. Note that the building masks are taken from the xBD dataset and not predicted by the model, they are used only for visualization purposes.