17th World Conference on Earthquake Engineering, 17WCEE
Sendai, Japan - September 13th to 18th 2020

MACHINE LEARNING APPROACH FOR EARTHQUAKE DAMAGE DETECTION FROM SPATIO-TEMPORAL REMOTE SENSING DATA

T. Miyamoto[1]

[1] *Associate Professor, University of Yamanashi, tmiyamoto@yamanashi.ac.jp*

Abstract

During large earthquakes, it is important to quickly determine the damage distribution in housing structures for disaster prevention measures. Currently, the acquisition of information is time consuming as it is only done manually by local public organizations. Therefore, a means to gather information promptly and objectively is required. As an effective tool for detecting disaster damage analysis of satellite imageries is widely studied. However, it is still difficult to determine building damages due to limitations in the resolution of satellite imagery. In this study, a system to detect building damage from a set of multi-temporal satellite imageries was developed by applying a recent machine learning approach. The two basic ideas to realize the system are as follows: First, using the information on the positions and shapes of structures stored in a GIS database, the photographic scope of each residence in wide-area photographic imagery was identified and small photographic fragments at individual structure level were extracted. Second, using a classifier, which determines whether the individual fragments depict collapsed structures or not, the damage in residential structures in the affected area was assessed. In this study, the effectiveness of the following two concepts for improving the performance of the classifier was evaluated. Using images of terrain periodically captured by satellites, images of the affected area for the two cycles (i.e., before and immediately after the earthquake) were fed simultaneously into the classifier to improve classification performance. The spatio-temporal convolution, which is considered a generalized method of image subtraction, was found the most effective. Then, rankSVM, a recently proposed machine learning model for handling imbalanced disaster-related data and the problems the pose to classification, was applied.

Keywords: machine learning, earthquake damage detection, spatio-temporal analysis, class imbalance

1. Introduction

Obtaining accurate information on the damage situation when natural disasters occur is the first priority for an appropriate response. Typically, information-gathering during emergency relies on the on-the-ground assessments of staff working in public organizations and reports from citizens. However, these measures designed by humans are highly time-consuming and gathering prompt information becomes more difficult when the scale of the disaster is large. Therefore, a faster and more objective method of gathering damage information is required.

One promising method for gathering damage information is the use of remote sensing technologies. Specifically, satellites can be used as platforms for optical sensors, which gather precise imagery, and for synthetic aperture radar, which can record terrain regardless of weather and daylight. By their nature, these technologies are well suited to gather damage information from natural disasters in terms of coverage area, speed, and weather tolerance among others. In addition, work on adapting these forms of information for disaster response is underway.

After an earthquake, it is important to assess building collapses at the micro level of individual residential structures to save lives. However, detecting such structural damage requires a spatial resolution beyond the limits of current satellite sensing data, which imprecisely detect damage; hence, making more precise detection is an important research focus.

Thus, this study proposed and developed a technology for assessing damage condition at residential level using wide-area photographic images from satellite sensors and by applying recent machine learning techniques. The proposed technology comprises two processes. First, the photographic scope of each residence in wide-area photographic images is identified using the information on the positions and shapes of structures stored in a GIS database. Then, small photographic fragments of each individual structure are extracted. Second, using a classifier, which determines whether these individual fragments depict collapsed structures or not, the damage in residential structures in the affected area is assessed.

In this study, the effectiveness of the following two concepts for improving the performance of the classifier was tested. First, using images of terrain periodically captured by satellites, images of the affected area for the two cycles (i.e., before and immediately after an earthquake) were simultaneously fed into the classifier to improve classification performance. Here, methods for extracting features from images before and after an earthquake were compared. Using spatio-temporal convolutional layers, which is considered a generalized method for image subtraction, was found most effective.

Additionally, the data related to disaster damage are typically imbalanced, in which acquiring data on damaged-seeming class is more difficult than on undamaged-seeming class. Thus, rankSVM, a recently proposed machine learning model that handles imbalanced data and the problems the data pose to classification, was applied. RankSVM does not pass through processes that distort the nature of data distribution, such as underpredicting or adjusting cost functions, for it can learn directly from imbalanced training data. Furthermore, it maximizes RoC-AUC, which is one of the indices for assessing classifier performance. It has important qualities as a classifier for imbalanced data related to the extent of damage. In this study, rankSVM was used to replace the output layer of a deep learning classifier. Its effects on classifier performance, particularly from the perspective of improving learning performance based on imbalanced data, was tested.

To verify the above two ideas, a dataset based on the satellite photography from the affected areas during the 2016 Kumamoto earthquake was constructed. The model was evaluated through a cross validation test. From the numerical experiment, it was found that spatio-temporal convolution of multi temporal images was highly effective and that learning from imbalanced training data that were not undersampled using rankSVM yielded a slight improvement in the classifier performance.

2. Related work

A comprehensive review by Dong and Shan [1] showed detection of building damage induced by an earthquake via remote sensing and acquisition of data through various sensors, such as the optical, SAR, and Lidar. Many studies have focused on the changes in the data observed before and after the disaster. Tong et al. [2] and Matsuoka and Yamazaki [3] proposed methods for determining structure collapse based on threshold index. The former calculated the difference in height of each structure derived from DEM before and after the earthquake, while the latter calculated the difference in the backscattering coefficient of the SAR radar. Further, recent studies have utilized machine learning models in detecting earthquake damage. Mansouri and Hamednia [4] applied a SVM model based on the features of an optical image difference before and after the disaster. Bai et al. [5] applied the K-nearest neighbor algorithm to the difference value of SAR images before and after the earthquake.

Satellite sensing for detection of an earthquake damage to a structure and techniques to classify the extent of damage based on pattern recognition models that determine the differences in the features of the images before and after the earthquake have often been adopted. It has been known that pattern recognition technology has drastically improved its performance in recent years with the progress of deep learning [6]. The convolutional neural network, whose effect has been demonstrated by Krizhevsky et al. [7], has improved the performance of classifiers by automatically constructing filters that extract features from an image via back propagation. Recently, many cases can be seen where CNN filters are extended in three dimensions (3D) and are used for the extraction of features from spatio-temporal data. Tran et al. [8] discussed the theoretical differences and advantages of 3D CNN with cases where a group of images are stacked in a feature dimension to which a 2-dimensional filter is applied. Ji et al. [9] applied a 3D CNN to human action recognition based on time-series images and demonstrated its effectiveness. This study tried to apply this recent machine learning technique to satellite sensing imagery with an aim of improving task processing capability for earthquake damage detection.

Data sets related to natural disasters generally have class imbalance problem. This is due to the lower incidence of natural disasters and damage and the higher difficulty in obtaining data corresponding to damage classes than to non-damage classes, including normal state. The conventional methods in dealing with imbalanced data on a machine learning technique include adjustment of data number via under sampling, for example, by discarding part of the data from a class with higher number of data or by adjusting the penalty in the cost function. The conventional approaches have been summarized by He and Garcia [10]. As these techniques may induce loss of important information and over fitting, different approaches have been proposed in recently. Wallace et al. [11] demonstrated the effectiveness of combining under sampling with an ensemble learning model. Cruz et al. [12] argued the effectiveness of applying a pairwise learning model to classification tasks from imbalanced data via learning to rank. The pairwise learning models use pairs of positive and negative training data; thus, no learning biases arise for a particular class. It also has the benefit of being able to learn imbalanced data without distorting the distribution characteristics of data by adjusting the data count. Among pairwise learning models, rankSVM [13] was the one used in this study. As explained later in the paper, rankSVM has the necessary property of performing classification tasks corresponding to anomaly detection, including detection of disaster damage.

3. Methods

3.1 Overview of the proposed scheme

The proposed earthquake damage detection system consists of processes shown in Fig. 1. First, information on the position and shape of residential structures is extracted from a GIS database of regions affected by an earthquake. Then, based on the information, the pixel position at which the concerned structure shown in the broad satellite imagery is identified and a small image segment corresponding to the structure is isolated. This operation is applied to all residential structures in the target region to create a small image for each structure. Then a classifier is applied to this small image to determine the collapse or non-collapse structure in the image, so as to determine the damage in each structure in the affected region.

Through this process, information extraction from a GIS database and identification of the location of the structures in the satellite imagery can be performed with ease as the location information is attached to each pixel in the satellite image. When creating a dataset used for an assessment experiment, which is described in section __ below, the validity of the program codes is manually confirmed. This paper presents the development of a classifier to be applied to small images.

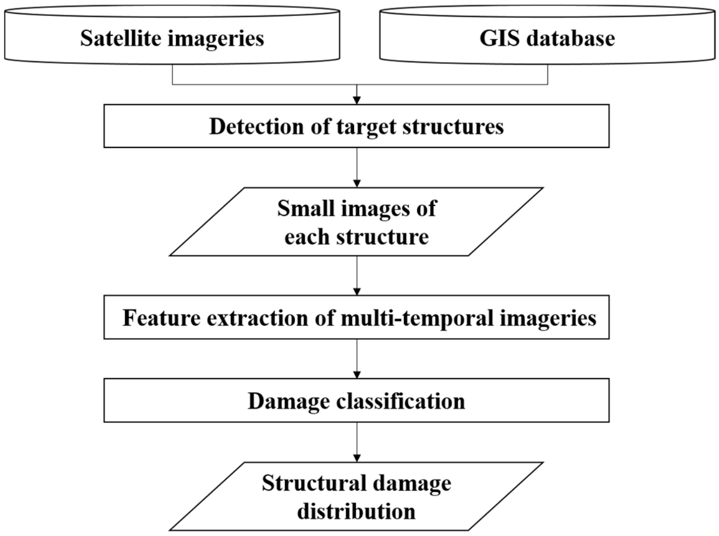

Fig. 1 – Flowchart of the proposed scheme

3.2 Feature extraction from multi-temporal satellite imagery

Since satellites periodically capture the surface of the earth, imageries right before the disaster, as well as imageries after the disaster can be obtained. Previous studies have reported that comparing the information obtained before and after a disaster can improve the accuracy of damage detection.

Most researches have used difference values for images before and after the disaster. However, the difference value is just one of the many characteristics that can be obtained from multi-temporal images. This suggests that the use of more features can potentially improve the accuracy of damage detection. Thus, in this study, a three-dimensional convolutional neural network (3D CNN) was applied as a classifier to solve classification tasks from a group of temporal images. The convolution layer in 3D CNN is an extension of the 2D filter in a normal CNN. In this study, a filter with one temporal dimension and 2 spatial dimensions was used. The 3D CNN automatically configures the values of filters through back propagation and extracts data features in temporal or spatial direction that is effective in performing the task.

This study used a deep learning model, which combined the feature extraction layers of the 3D CNN with the fully connect layers as a classifier to classify quake-induced damage from the multi-temporal images.

3.3 Learning from imbalanced disaster data

Generally, there is less data that corresponds to times of disaster than to normal or non-damage times. Therefore, the issue of class imbalance is one of the important problems that must be addressed in detecting disaster damage. As mentioned later in this paper, among images used in this research for the cross validation test, there was a greater number of non-damaged structures than damaged structures. Thus, this data set had a class imbalance issue.

In this study, rankSVM was adopted as a classifier to address imbalanced data. The rankSVM is originally one of the machine learning models that are used for learning to rank. Recently, its application for binary classification of class imbalance data has been proposed.

If positive class data $\boldsymbol{P}=\{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_n\}$ and negative class data $\boldsymbol{N}=\{\boldsymbol{x}'_1, \boldsymbol{x}'_2, \ldots, \boldsymbol{x}'_m\}$ are given, binary classification problems can be solved by an evaluation function f, where, for arbitrary $\boldsymbol{x}_i$ and $\boldsymbol{x}'_j$, the following formula holds.

$$f(\boldsymbol{x}_i) > f(\boldsymbol{x}'_j) \tag{1}$$

The rankSVM constructs the evaluation function f with feature conversion mapping φ corresponding to an arbitrary kernel and performs linear conversion as follows.

$$f(\boldsymbol{x}) = \boldsymbol{w} \cdot \varphi(\boldsymbol{x}) \tag{2}$$

The optimal value of $\boldsymbol{w}$ is obtained by solving the following optimization problem [14].

$$\boldsymbol{w}_{\text{opt}} = \underset{\boldsymbol{w}}{\operatorname{argmin}}\, C \sum_{(i,j)} l\left(\boldsymbol{w} \cdot \varphi(\boldsymbol{x}_i) - \boldsymbol{w} \cdot \varphi\left(\boldsymbol{x}'_j\right)\right) + \frac{1}{2}\|\boldsymbol{w}^2\| \tag{3}$$

Using the solution of the equation for the evaluation function, the rankSVM acquires the following two important properties of a classifier for imbalanced data.

Firstly, in the equation of the optimization problem, the positive data $\boldsymbol{x}_i$ and the negative data $\boldsymbol{x}'_j$ are always used in a pair as a training data sample. Therefore, despite the difference in the number of data between two classes, a class with a larger data count would not appear more in the training data sample than the other class and the learning bias associated with the imbalance in data count would not occur. Thus, using the rankSVM, imbalance data can be used directly without resampling or adjusting the cost function and without undermining the properties of the original data distribution.

Secondly, the first term on the right side of the optimization problem equation represents a loss due to pair (i,j) that does not satisfy the condition $f(\boldsymbol{x}_i) > f(\boldsymbol{x}'_j)$. Thus, the minimization of this term works in a way that maximizes the number of (i,j) pairs that satisfy the equation. The ROC-AUC, which is one of the indices used to evaluate performance of classifiers, can be represented as follows using the evaluation function f [15].

$$\text{ROC-AUC} = \frac{|(i,j);\, f(\boldsymbol{x}_i) > f(\boldsymbol{x}'_j)|}{|\boldsymbol{P}| \times |\boldsymbol{N}|}, \tag{4}$$

where $|\boldsymbol{P}|$ and $|\boldsymbol{N}|$ represent the positive and negative data count, respectively and $|(i,j)\,;f(\boldsymbol{x}_i) > f(\boldsymbol{x}'_j)|$ represents the number of pairs of a positive and a negative data that satisfies $f(\boldsymbol{x}_i) > f(\boldsymbol{x}'_j)$. Taking all these considerations together, the evaluation function f in the rankSVM is learned to maximize the ROC-AUC [14]. If the ROC-AUC value is high, it means that a higher true positive rate has been achieved in a lower false positive rate. This is an important property as a classifier for anomaly detection, such as in disaster damage detection.

3.4 Experimental study

3.4.1 Data set

The performance of the machine learning model described above was verified using the results from a field research conducted in Mashiki town, which was one of the affected areas of the 2016 Kumamoto earthquake in Japan. A dataset of satellite images of the affected residences, which were composed of 326 collapsed buildings and 648 unaffected buildings was constructed. In this study, more focus was given on the data obtained through optical sensors. From a set of pre- and post- disaster images captured at two time points [the first on December 15, 2015 and the second on April 29, 2016 (Fig. 2)] by the optical satellites Spot 6 & 7, a small image segment for each residence was obtained. The data resolution was 1.5 m/pixel, and each image was resized to 40×40 pixels in order to adjust the size, which varied according to the footprint of the residence.

(a) Pre-disaster image (December 15, 2015) (b) Post-disaster image (April 29, 2016)

Fig. 2 – Pre- and post- earthquake imagery of Mashiki town in Japan taken by the optical satellite SPOT 6 & 7

3.5 Evaluation on the effectiveness of 3D CNN

To test the validity of using a 3D CNN as a means to extract features for the classification of disaster damage from multi-temporal images, various input tensor forms including difference values of images and the corresponding deep learning models were developed (Table 1) and their classification performances were compared.

Table 1 – Overview of the test cases for comparing the input tensor forms and the corresponding convolutional neural network models

No	Input data	Input tensor form		Dimension of the tensor form	Deep learning model
0	Post-earthquake image	2D image		$w \times h \times 3$	2D CNN
1	Pre- and post- disaster images	Image difference values	−	$w \times h \times 3$	2D CNN
2	Pre- and post- disaster images	Stacking images in feature dimension	+ 1-3ch 4-6ch	$w \times h \times 6$	2D CNN
3	Pre- and post- disaster images	Stacking images in time dimension	+ t_1 t_2	$w \times h \times t \times 3$	3D CNN

In case No. 0, an input tensor was developed using a post-disaster image only. The tensor had three dimensions: width (w), height (h), and RGB (3ch) and a normal 2D CNN was used.

Case No.1 represented a case where an input tensor was developed by calculating the image difference values of pre- and post-earthquake images. Here, the number of dimensions of the tensor, as well as the CNN model, was similar in No. 0; however, it differed from No. 0 as it used pre-disaster images.

In case No. 2, a $w \times h \times$6ch input tensor was constructed by stacking a post-earthquake image as 4 to 6ch data, following a 3ch pre-earthquake image. Here, the computation process of the CNN was identical to that of the normal CNN except for 6ch for the number of features in layers. However, contrary to case No. 2, the data were directly input without calculating the difference between the images before and after the earthquake. Hence, the input tensor was expected to be more informative.

Case No. 3 was one where an input tensor was constructed as a $w \times h \times t \times 3$ch data by stacking pre and post-earthquake images in the temporal dimension. The computation process of the CNN was 3D. It was expected that, through back propagation, feature extraction filters that capture more detailed features in the spatial and temporal direction could be formed.

By comparing the performances of the model under these four cases, the input data form and deep learning model that are suitable for the detection of quake-induced damage can be discussed. Note that, for comparison, the number of layers in the model was equalized, and the number of internal parameters was almost equalized in each case. To illustrate this, the model architecture of case 3 is shown in Fig. 3.

model architecture	shape of filters ($w \times h \times t \times ch$)	output tensor dimensions ($w \times h \times t \times ch$)
Convolution3D	5×5×2×5	40×40×2×5
Convolution3D	5×5×1×10	40×40×2×10
Maxpooling3D		20×20×2×10
Convolution3D	5×5×1×20	20×20×2×20
Convolution3D	5×5×1×20	20×20×2×20
Maxpooling3D		10×10×2×20
Convolution3D	5×5×1×20	10×10×2×20
Maxpooling3D		5×5×2×20
Fully connected	256	1×1×1×256
Fully connected	10	1×1×1×10
Fully connected	1	1×1×1×1

Fig. 3 – Model architecture of 3D CNN for case 3

The number of data used in the cross validation test is shown in Table 2. In this comparative test, images in the non-damage class were randomly under-sampled from the original dataset in order to equalize the data count between the two classes.

Table 2 – Number of images used for damage detection
(Data augmentation is applied to the training data.)

	Non-Damaged	Damaged
Training Data	3,880	1,304
Test Data	163	163

3.6 Evaluation on the effectiveness of rankSVM

The classification performance of rankSVM for an imbalanced disaster data set was also verified as follows.

From a pre-trained 3D CNN model used in the above test, fully connected layers were replaced by an SVM classifier (Fig. 4), which in turn learned from an imbalanced training data as shown in the Table 2. Then the classifier was tested for its performance against the test data. Three kinds of SVMs as classifiers were tested: normal SVM, one class SVM, and rankSVM. The hyper parameter values for each SVM as well as the type of kernel was determined via grid search. Moreover, a case in which the number of training data was equalized by under sampling was compared with a case where imbalanced training data were straightforwardly used for learning. For normal SVM and one class SVM, models implemented in scikit-learn library were used. RankSVM was implemented from scratch in this study following Pegasos algorithm [16].

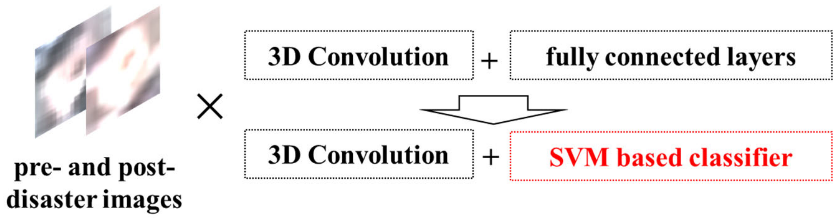

Fig. 4 – Classifier by replacing fully connect layer with SVM

4. Results and discussion

4.1 Comparison between input tensor forms

The result of the comparison test for the evaluation of 3D CNN is summarized in Table 3. The performance of the models was evaluated in terms of accuracy, precision, recall, F-measure, and ROC-AUC. Compared to case 0, cases 1-3, where both pre- and post-earthquake images generally indicate better performance, show the effectiveness of multi-temporal satellite imagery for disaster damage detection. The 3D CNN model in case 3 has the best classification performance among the models compared. As the 3D convolution filters can extract more general features from spatio-temporal data including image difference values, it is considered that deep learning models acquire features that contribute to earthquake damage detection through the learning process.

Table 3 – Comparison of the model performance

No	Accuracy	Precision	Recall	F-measure	ROC-AUC
0	0.641	0.883	0.595	0.711	0.715
1	0.687	0.564	0.748	0.643	0.785
2	0.702	0.669	0.717	0.692	0.779
3	0.761	0.669	0.820	0.736	0.819

As indicated by the high ROC-AUC value, the model in case 3 showed good performance as indicated by the ROC curve (Fig. 5). Comparing to other models, higher true positive rates can be achieved under a certain false positive rate by the classifier model in case 3. In disaster damage detection, it is important, in the context of safe decision making, not to overlook severe damage, even to a point of overestimation. The high ROC-AUC value of 3D CNN is an essential property of a damage detection model.

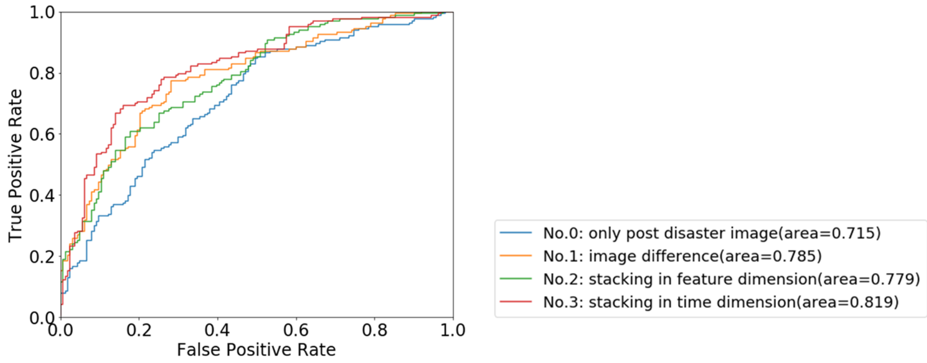

Fig. 5 - Comparison of ROC curves

4.2 Learning from imbalanced data

Table 4 shows the classification performance of SVM classifiers using learning balanced or imbalanced training data. The performance of fully connect layers is also shown, which presents the results before replacing the classifier layers in Fig. 4. As a result of grid search, the Gaussian kernel was applied in all SVM model.

Table 4 – Performance of SVM classifiers

Classifier	Training data balance	Accuracy	F-measure	ROC-AUC
Fully connected layer	Under sampling	0.761	0.736	0.819
Normal SVM	Under sampling	0.739	0.738	0.739
	Imbalanced	0.622	0.667	0.623
One class SVM	Under sampling	0.647	0.677	0.647
	Imbalanced	0.638	0.673	0.638
RankSVM	Under sampling	0.770	0.773	0.820
	Imbalanced	0.770	0.778	0.821

The normal SVM classifier did not perform better than the fully connected layer. Furthermore, the performance of SVM with imbalanced training data was lower than that with balanced data, which is a typical problem of imbalanced data. The decrease in the performance of one class SVM was not remarkable and the performances were almost similar between the different training data balance. In the training process of one class SVM, only the data for the non-damaged class, whose data count was larger than the other class, were utilized. Thus, there was no need to discard the non-damaged class data to balance the data count. Instead, the data for the damaged class were excludedfrom the learning process. The results showed that the performance of one class SVM was not better than the fully connected layer in this case. Note that one class SVM is usually an effective tool for anomaly detection. However, in this case, the features of normal data may have been diverse because the data consisted of various types of non-damaged housing structures. Learning the classification surface of one class SVM had become difficult probably due this diversity.

However, using the rankSVM as the classifier layers resulted in a slightly better performance. Moreover, the performance of rankSVM with imbalanced training data improved slightly than that with under sampling case. Though the improvement was small, the result indicated that the classification performance was possibly enhanced by increasing only the number of non-damaged data, which can be easily obtained. The high learning potential from imbalanced training data, which rankSVM possesses is an important property of a classifier for disaster damage detection.

Figure 6 shows the inference result obtained by the proposed classifier with the combination of 3D CNN and rankSVM. Note that both the inference results from 648 training data and that from 326 test data were included in this figure. Compared to the ground truth data, the inference results reproduced better the extent of damage in each region.

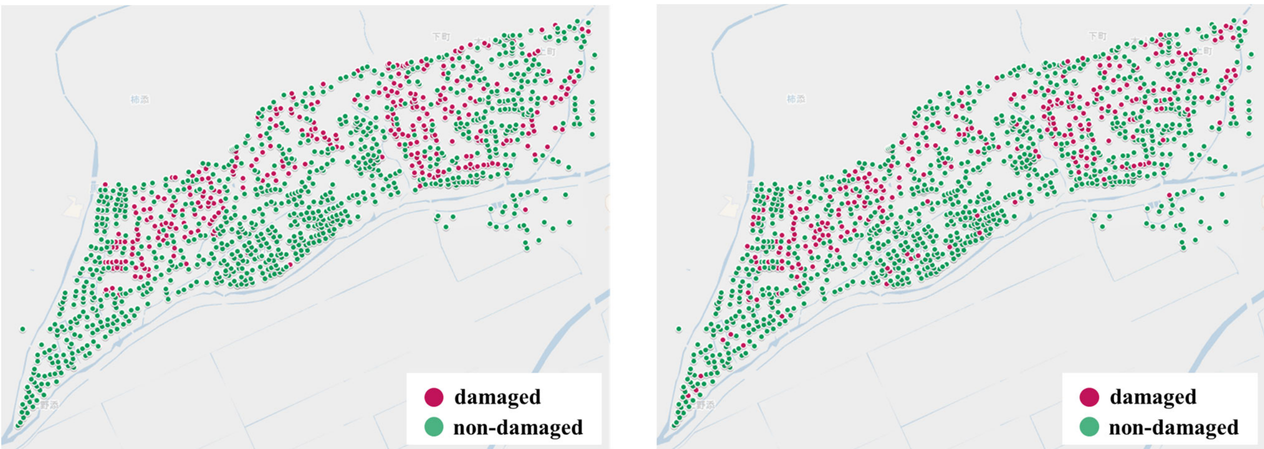

(a) Ground truth (b) Inference

Fig. 6 – Result of earthquake damage detection by the proposed scheme (©2020 Google)

5. Conclusion

This study proposed a scheme for earthquake damage detection using satellite imagery and recent machine learning techniques. Using the properties of time series images and class imbalance in satellite imagery, this study verified the effectiveness of 3D CNN and rankSVM. Experimental results show that 3D CNN performed better than with the other schemes for feature extraction from multi-temporal images. Moreover, although the improvement in the performance obtained from replacing the classifier layers by rankSVM was slight, it indicated a possibility that classification performance can be enhanced by increasing only the number of non-damaged data, which can be easily obtained. The results show that the proposed method is suitable for the properties of satellite imagery.

6. References

[1] Dong L, Shan J (2013): A comprehensive review of earthquake-induced building damage detection with remote sensing techniques. *ISPRS Journal of Photogrammetry and Remote Sensing*, **84**, 85-89.

[2] Tong X, Hong Z, Liu S, Zhang X, Xie H, Li Z, Yang S, Wang W, Bao F (2012): Building-damage detection using pre- and post-seismic high-resolution satellite stereo imagery: A case study of the May 2008 Wenchuan earthquake. *ISPRS Journal of Photogrammetry and Remote Sensing*, **68**, 13-27.

[3] Matsuoka M, Yamazaki F (2004): Use of satellite SAR intensity imagery for detecting building areas damaged due to earthquakes. *Earthquake Spectra*, **20** (3), 975-994.

[4] Mansouri B, Hamednia Y (2015): A soft computing method for damage mapping using VHR optical satellite imagery. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, **8** (10), 4935-4941.

[5] Bai Y, Adriano B, Mas E, Koshimura S (2017): Machine learning based building damage mapping from the ALOS-2/PALSAR-2 SAR imagery: case study of 2016 Kumamoto Earthquake. *Journal of Disaster Research*, **12**, 646-655.

[6] LeCun Y, Bengio Y, Hinton G (2015): Deep learning. *Nature*, **521**, 436-444.

[7] Krizhevsky A, Sutskever I, Hinton GE (2012): ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NIPS)*, 1106-1114.

[8] Tran D, Bourdev L, Fergus R, Torresani L, Paluri M (2015): Learning spatiotemporal features with 3D convolutional networks. *IEEE International Conference on Computer Vision (ICCV)*, 4489-4497.

[9] Ji S, Xu W, Yang M, Yu K (2013): 3D convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **35** (1), 221-231.

[10] He H, Garcia EA (2009): Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, **21** (9), 1263-1284.

[11] Wallace BC, Small K, Brodley CE, Trikalinos TA (2011): Class imbalance, redux. *11th IEEE International Conference on Data Mining*, 754-763.

[12] Cruz R, Fernandes K, Cardoso JS, Costa JFP (2016): Tackling class imbalance with ranking. *International Joint Conference on Neural Networks*, 2182-2187.

[13] Joachims T. (2002): Optimizing search engines using clickthrough data. *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 133-142.

[14] Chapelle O, Keerthi SS (2010): Efficient algorithms for ranking with SVMs. *Information Retrieval*, **13** (3), 201-215.

[15] Joachims T. (2005): A support vector method for multivariate performance measures. *Proceedings of the 22nd International Conference on Machine Learning*, 377-384.

[16] Shalev-Shwartz S, Singer Y, Srebro N, Cotter A (2011): Pegasos: primal estimated sub-gradient solver for SVM. *Mathematical Programming*, **127** (1), 3-30.

