

Rapid, Automated Post-Event Image Classification and Documentation

Chul Min Yeum¹, Shirley J. Dyke¹, Bedrich Benes², Thomas Hacker³, Julio Ramirez¹,
Alana Lund¹, and Santiago Pujol¹

¹ Lyles School of Civil Engineering, Purdue University, West Lafayette, IN, 47907, United States.
{cyeum, sdyke, ramirez, alund15, spujol}@purdue.edu

² Computer Graphics Technology, Purdue University, West Lafayette, IN, 47907, United States.
bbenes@purdue.edu

³ Computer and Information Technology, Purdue University, West Lafayette, IN, 47907, United States.
tjhacker@purdue.edu

Abstract. Reconnaissance teams collect large volumes of perishable data after a natural disaster that are related to the condition of the buildings and other infrastructure. Each event is an opportunity to evaluate the performance of our structures under circumstances that cannot entirely be reproduced in the laboratory or numerical simulation. In the field, engineers typically prefer to record such information through images. For each building, images readily document the visual appearance of damage to structures and their components. Each team follows a similar procedure that includes taking views of the structure, from both outside and inside, using various distances and angles. During this process, engineers frequently incorporate metadata in the form of images, such as structural drawings, GPS devices, watches, and even measurements (e.g. an image of a structural column with a measuring tape). Large quantities of images with a wide variety of content are collected within a short period, and their timely organization and documentation are important. Engineers need to generate accurate and rich descriptions of the images before the details are forgotten.

To distil such information in an efficient and rapid manner, we developed an automated approach that uses computer vision techniques for classifying and organizing the large volume of images collected from each building. Deep convolutional neural network (CNN) algorithms are successfully implemented to extract robust features of key visual contents in the images. This capability is demonstrated using data collected by various reconnaissance teams from buildings damaged during past earthquakes. A schema is developed based on the real needs of field teams examining buildings, with several different categories and information defined and used for annotating the images, thus supporting organization of the data that parallels the procedure followed by the engineer in the field. A significant volume of these images from past earthquakes is used to train robust classifiers that can automatically classify the images. These are then used to automatically generate individual reports for buildings that were damaged in past earthquakes.

Keywords: Post-disaster evaluation, Convolutional neural networks, Image classification, Building reconnaissance.

1. INTRODUCTION

After a disaster, post-event building reconnaissance teams are charged with collecting perishable data to be used for learning from disasters before that data is destroyed. The primary functions of these teams are to collect readily available, perishable data to enable scientific research intended to: (1) learn as much as possible about the nature of the event and extent of the consequences; (2) identify potential gaps in existing research or in the practical application of scientific, economic, engineering or policy knowledge; and (3) make recommendations regarding the need for further investigations, and/or changes to codes, standards and design guidelines. Damaged structures and their components provide critical information regarding performance during the event, and lessons learned from structures that do not experience damage are just as important [Yeum, 2016].

In a typical mission, a group of data collectors is dispatched to a region where an event has taken place. In a well-organized team, information about the local construction, severity of the event, as well as maps of the region are often made available in advance so that planning can take place. The larger group is divided into small teams with at least one or more experienced structural engineering evaluator on each team. Each team visits 4-5 buildings a day, collecting images at each site and taking measurements from each building. The teams may follow the procedures outlined in established guidelines (e.g. ATC-20 (earthquakes) and ATC-45 (windstorms and floods)) for this process, which is intended for rapid structural evaluation after events (although these teams are not directly rating these buildings) (www.atccouncil.org). An enormous amount of visual data (images and videos) can be generated in just a few days. Each evening teams return to the base to discuss the findings and to review plans regarding where to spend time and effort collecting data on the next day.

In the field, engineers often record new findings and observations using photographs. Meaningful scenes are captured from unique viewpoints and locations that can be helpful to recall and confirm them in the future. The subjects of these images include the visual appearances of the damaged/undamaged buildings and their components, but will also encompass several types of metadata that are needed in the documentation (e.g. drawings, GPS navigator). Since photographs are sequentially collected throughout each building, they also contain valuable temporal and spatial contexts. For example, an image of structural components taken together with a measuring tape will include the appearance of the target component as well as the measurement value. An observation leading to an image taken on a higher floor of the structure may trigger the need for further images on the ground floor. Thus, engineers can estimate the relative location of the components by comparing images that are collected according to a given sequence, also reducing the risk of human errors under strenuous field conditions.

A major challenge of dealing with such images in the field is that they must be organized and documented in a timely manner. For example, reconnaissance teams must make decisions regarding where to next allocate teams and resources in just minutes to hours. Additionally, some images may require an explanation of the purpose and intention behind the image, which must be provided by the engineer who collected the images. Thus, immediate documentation is critical before the important details are forgotten. With a large and varied amount photographic data, manual sorting of these large volumes of images is tedious and time-consuming. Reconnaissance teams can save tremendous time and effort if automation is available for classifying those data and displaying them in a favorable way to streamline this process.

In this study, we develop the enabling technology, and an associated schema, to automatically extract and analyze the metadata and visual contents of such data sets to classify and organize them. Deep convolutional neural network (CNN) algorithms are implemented to extract distinct features for categorizing the images according to a pre-determined set of appropriate classes. Using our large-scale, real-world image data set, robust image classifiers are trained for each class in the schema. In addition, additional information including date, time, or GPS information is extracted from exchangeable image file format (hereafter, EXIF) for each

image. Extracted information is systemically integrated and organized to enable rapid generation of reports according to a suitable digital format so that the engineers can easily access and describe the images. This report generation capability is intended to provide rapid decision support for post-event reconnaissance field investigations, as well as documentation for future data reuse. With this capability, engineers can browse through the report for a given building and rapidly find and examine particular images to identify an important image-of-interest. To demonstrate the technology and its capabilities, we first build a large-scale, ground-truth image database to train the classifiers of the classes in the schema that are used for generating the report. This database contains real images collected during past earthquake reconnaissance missions from around the world. According to the schema developed to support earthquake report generation, multiclass and binary classifications are combined to support the class hierarchy by arranging distinctive image classes in each training scenario. Using these trained classifiers, sample reports are automatically generated from large sets of images collected from buildings damaged in the Taiwan and Ecuador earthquakes in 2016 [Purdue University and NCREE, 2016; Sim *et al.*, 2016].

The merit of this study is that we provide a practical and feasible solution to support actual field engineers in tasks that could not be automated without this capability. We do not create artificial images for either training classifiers or evaluating them. Instead, we use real-world sets of complex and unstructured images using an unprecedented number of real-world images. The images were collected from several previous earthquake events, and contain hundreds of buildings from many regions and with various appearances. Thus, the developed solution is unbiased, trustworthy and practical, and it can be immediately deployed during actual reconnaissance missions in the future.

The remainder of this paper is structured as follows. Section 2 begins with the overview of the approach. Image schema used in this study are illustrated in Section 3. In Section 4, the capability is demonstrated using real-world images collected from past post-earthquake reconnaissance missions. Section 5 includes the summary and conclusions.

2. OVERVIEW OF THE APPROACH

The concept behind the approach is illustrated in Figure 1. The main goal is to directly support the planning of subsequent data collection in the region by rapidly gathering and providing some organization and structure to the data as they are collected. Visual content and metadata are maximally extracted for classifying and cataloging the data collected. First, a large volume of images is collected from several buildings by teams of field engineers. In addition to the images that provide the visual appearance of damage, these image collections frequently incorporate metadata in the form of images, such as photographs of drawings or measurements. Depending on the size and/or damage condition of the building, several hundred images are often collected within a short time. Next, pre-trained classifiers are used on the entire data collection to automatically analyze and extract the visual contents and metadata, turning the images into information of use to the team. The schema to be used in organizing the images are designed in advance based on the needs of the given application. In this study, the classes we used were designed according to two groups: building and buildings components (e.g. inside or outside the building) and metadata recording (e.g. GPS, drawing). Each image is then annotated with the results of the classification process, at each level in the hierarchy. Finally, the extracted and classified information is organized into a report. Although many possible formats could be used for various purposes, we have developed a format that will directly support the earthquake reconnaissance teams. The report begins with a small set of representative overall images and a description of the metadata information about the site and image collection. Beginning the report with this image will quickly identify the subject building from the several that were investigated during the mission. Then, images are shown in the classified category in order according to the time they were collected.

Two sample reports are shown in Figure 5. Rather than looking at a sequence of mixed and unstructured individual images, the team can efficiently and rapidly explore the images using the report. Note that the report is automatically generated from a set of images and no manual process is used, although users can freely add comments to the images or the report.

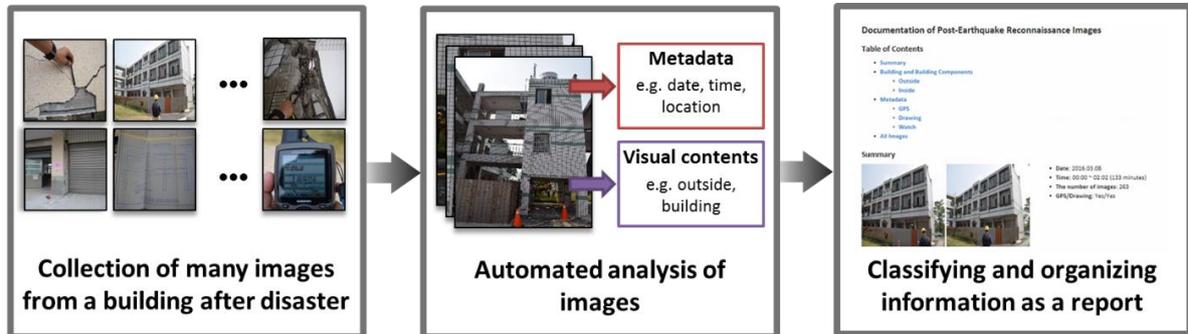


Figure 1. Overview of the approach: A large volume of images collected during post-event reconnaissance missions (left) are automatically turned into a well-organized report by analyzing their visual contents and metadata (middle) that is used for an automatic creation of mission reports (right).

3. DESIGN OF THE SCHEMA

To support the application, a hierarchical structure was developed based on discussions with field engineers. Classes must be useful for the application, but classification can only proceed if they are also visually distinguishable. The hierarchy also enables efficiency and accuracy because visual contents in images are classified in order. For instance, classes such as outside building images and inside building images should only be applied to images that are already classified as being related to buildings and their components.

To successfully classify the images to the categories, a clear definition is needed for each category. Definitions are essential because they guide human annotators to establish a consistent and meaningful ground truth data set for training. Additionally, the classification results will be more accurate when there are clear boundaries between the features in the images of different classes. We describe our definition for each category. These definitions are used for labeling our images by human annotators for the demonstration in Section 4. Sample images in each class are presented in Figure 2:

- **Building and building components** (hereafter, *BBC*): This category is applied to images that contain the visual appearance of a physical building, including either inside or outside images. Examples of such images are those that contain undamaged or damaged structural components, non-structural building and building components, close-ups of damaged areas, or building components either from the inside or outside. Images containing measurements of structural components or damage are included in this category. Images taken to record metadata, such as drawings or GPS devices, are not included in this category. Images in this category will be further classified (see below).
- **Overview** (hereafter, *OV*): Overview images are defined by those that represent the complete appearance of the building, usually from a distance. These images allow the users to estimate how many floors are present, observe the architectural type or structural type, etc. The entire side and/or front view of the building should be included. Partial obstruction may be included, but roughly 70 % of the building should be visible. Images containing the entire building façade, diagonal view of the building, or a complete side of the building are included in this category.

- **Building inside** (hereafter, *BIN*): This category is defined as images that have a general sense of inside space. For examples, these images might view a space surrounded by walls or a window, and they might contain interior building components or indoor space (e.g. room, basement, corridor).
- **Building outside** (hereafter, *BOUT*): This category has the opposite meaning to that of the building inside. The images show a space having no ceiling or surrounding walls.
- **Measurement** (hereafter, *MEAS*): Images in this category are to record the absolute or relative size of the building components (e.g. column, rebar spacing) or damage (e.g. crack, spalling). The size may be measured using a measuring tape, crack gauge or fixed size objects (e.g. pen, hand, note). The images are typically captured at a close distance to the target object so that the digit on the measurement tool is visible.
- **Watch** (hereafter, *WATCH*): Images that have an analog or digital timepiece are commonly collected. The purpose is to document the time associated with a particular stage in the data collection process, in case images are reformatted/resized and the EXIF information is lost or the date and time in cameras are not adjusted for the local time zone.
- **GPS** (hereafter, *GPS*): Some images are collected to show the values measuring GPS location on a GPS receiver. Again, these are also helpful in case the geotagged images are reformatted and the EXIF information is lost.
- **Document** (hereafter, *DOC*): This category is defined as an image of text and/or drawings. For example, drawings, sign or text documents are included in this category. Often digital records of the structural drawings for the subject buildings are not available. Thus, photographs are taken of the entire set of structural drawings to provide valuable reference information about the subject building.

Note that these categories are designed for this application, and can be expanded or modified according to the needs of a particular application and the quality of the images.

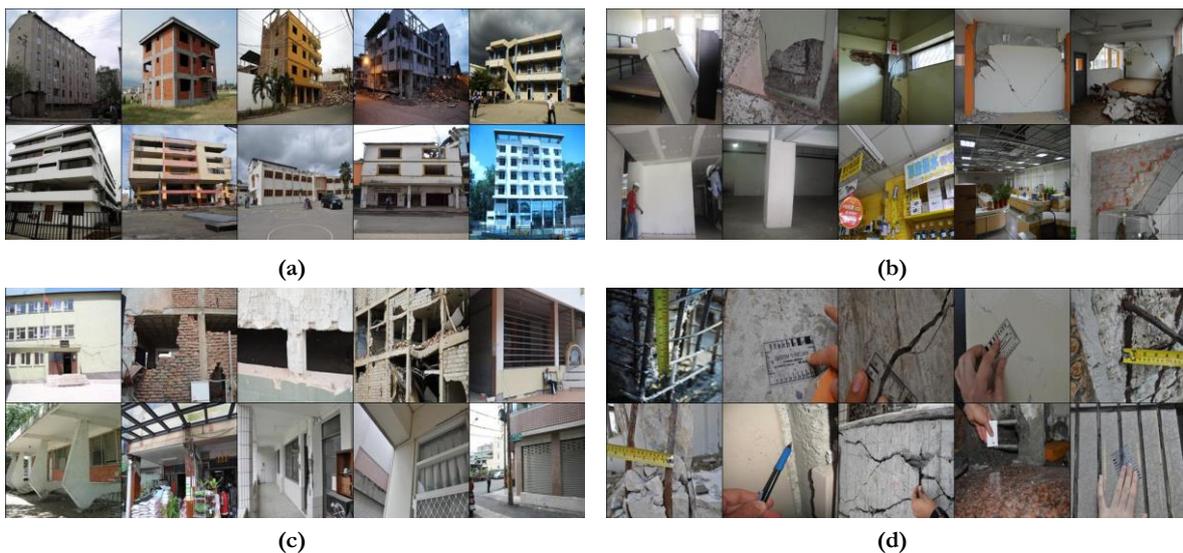




Figure 2. Sample images in each of the classes designed for this demonstration: (a) building overview, (b) building inside, (c) building outside, (d) measurement, (e) watch, (f) GPS, and (g) documents. Note that these labeled images are utilized for training classifiers to analyze the visual contents of the images. Here, all images are rescaled as square thumbnails for the arrangement.

4. EXPERIMENTAL VALIDATION

In this section, we describe the process used for the case study of earthquake images and demonstrate the accuracy of the approach on several complex, real-world, unstructured data sets. Similar capabilities for other applications could readily be implemented using a similar procedure and a well-defined schema.

4.1 GROUND-TRUTH LABELING USING OUR POST-DISASTER IMAGE DATABASE

We start with introducing our post-event reconnaissance image database [Yeum *et al.*, 2016a, 2016b]. We pulled a large subset of images from our database to train a classifier for each category as well as to demonstrate the capabilities of the approach. We have gathered an extensive collection of approximately 100,000 color images that were acquired by researchers and practitioners after past natural disasters. The database includes hurricane, tornado, and seismic events (e.g., from datacenterhub.org at Purdue University, disaster responders, Canterbury Earthquake Digital Archive, or the Earthquake Engineering Research Institute image collection) although for this demonstration we focus on earthquake images. Some sample collections included in our database are shown in Figure 3 [Sim *et al.*, 2015; Shah *et al.*, 2015; Sim *et al.*, 2016; Purdue University and NCREE, 2016]. Nearly all of the images we have preserved the original quality (resolution) as well as the basic metadata (e.g. date, time, and event), and a relatively small portion of the images contain GPS information or a picture of a GPS navigator. To date, the distribution of the types of disasters is earthquake (90%), hurricane (5%), tornado (4%), and others (1%). These images are collected from several different events such as earthquakes (e.g. Haiti in 2010, L'Aquila in 2009, Christchurch in 2011, Nepal in 2015, Taiwan in 2016), hurricanes (e.g. Florida in 2004, Texas in 2008), tornadoes (Florida in 2007, Greensburg in 2007) (datacenterhub.org). We will continue to collect images from such natural disasters to integrate into the database [Yeum, 2016].



Figure 3. Sample image collections from our post-disaster image database

Among the several collections of the images in our database, for this research, we utilize the images collected from earthquake reconnaissance missions in Düzce, Turkey in 1999, Bingöl, Turkey in 2003, Peru in 2007, Haiti in 2010, Nepal in 2015, Taiwan in 2016, and Ecuador in 2016 [Sim *et al.*, 2015; Shah *et al.*, 2015; Sim *et al.*, 2016; Purdue University and NCREE, 2016]. These images are available to the public at datacenterhub.org. A total of around 800 sets of building images are available for these events, and reconnaissance data for each building are documented, including measurements of structural components, building damage levels, and its images. Additionally, images that are not collected and organized in relation to a specific building at datacenterhub.org are also added to increase the number of images for training. We labeled these images into the developed classes in the schema.

For labeling the images that are to provide ground-truth training data, we design a web-based annotation tool to allow multiple annotators can work together to complete the task within a short time. This annotation modeled the tool that is used for a large-scale data [Yeum, 2016]. Users can rapidly annotate images using only a keystroke and can make comments if needed. Three annotators with training in civil engineering took part in the annotation of these images. Each labeled image is reviewed separately by at least two annotators. In the beginning, the annotators are given enough time to learn class categories with sample images and labeling. if there is a need for discussion of the class.

A total of 21,185 images were labeled using the classes designed in this study. The number of images in each class is presented in Table 1. *BBC* images include *BIN*, *BOUT*, *OV*, and *MEAS*. *BOUT* images include those from outside of the building as well as *OV* images. Here, only 103 *WATCH* images are in our database because some reconnaissance teams do not collect these images. Thus, we supplemented the training data for this class of images with a Google image search. A combination of keywords of “wrist” and “watch” is used for searching for these images. Using this approach, 207 images are added to the ground-truth datasets for *WATCH*. Samples of these images are marked as red in Figure 2(e). For the demonstration of report generation in Section 4.4, we select two sets of building images among our labeled dataset. These images are selected such that they are not used for training and testing classifiers in Section 4.2.

4.2 CONFIGURATION OF CONVOLUTIONAL NEURAL NETWORK

In this study, we implement a popular ImageNet CNN model called Alexnet (TorontoNet in Caffe), framed in MatConvNet library [Vedaldi and Lenc, 2014]. Alexnet exhibited superior implementation of CNNs in computer vision applications in the ImageNet image classification competition in 2012 and has been widely used for a benchmark test of a CNN model [Krizhevsky *et al.*, 2012]. The network architecture is presented in detail in the following reference [Russakovsky *et al.*, 2015]. Although some improved architectures have also been introduced and demonstrated after 2012, in this study, we choose Alexnet as an established method to implement this proof of concept of a simple and general CNN model.

For training a classifier of each category, labeled images are first transformed into inputs for the CNNs. Regardless of the category, all images are isotopically (preserving their aspect ratio) resized so that the shorter side of the images becomes 256 pixels, followed by cropping the center square region, to create 256 x 256 (pixels) square images. For data augmentation, 227 x 227 patches are randomly cropped from the 256 x 256 images in each epoch and random horizontal flipping is allowed. Additionally, random color shifting is applied to vary (jitter) intensities of RGB images.

We modify the last 1000-way softmax layer in the original implementation from the ImageNet competition to 4-way softmax for multiclass classification and logistic for binary classification [Krizhevsky *et al.*, 2012]. The layers are initialized as a Gaussian distribution with a zero mean and variance equal to 0.1. The hyperparameters are the same as those used in Alexnet [Krizhevsky *et al.*, 2012]. We trained our models using stochastic gradient descent with a batch size of 256 images, the momentum of 0.9, and weight decay of 0.0005. We train the network for 300 epochs, and the learning rate is logarithmically decreased from 0.01 to 0.0001 during training. To train the classifiers and test their capabilities, all labeled images are divided into 50%, 25%, and 25% for training, validating, and testing, respectively. A computer PC workstation having a Xeon E5-2620 CPU and NVidia Titan X with 12 GB video memory GPU is used for training and testing the algorithm. The MatConvNet library installed on Matlab 2017a is used for this study [Vedalidi and Lenc, 2014].

In our training dataset, the number of labeled images are highly unbalanced because the number of images intended to record metadata is much fewer than the number that are *BBC* images (*MEAS* images are categorized as *BBC* but their purposes are metadata recording). For example, the number of *BBC* images is around 50 times as many as *WATCH* images. If the training samples are assigned by randomly selecting these labeled images, the classifier would tend to overfit to the majority classes (herein, *BBC* and *DOC*). To tackle this issue, we assign the same number of images in each to a batch. For example, for 4-way multiclass classification, 128 images from each class are randomly selected if the size of the batch is 512. In such case, many repetitions of the images in the minority classes (e.g. *GPS* or *WATCH*) are permitted to fill the balance of the training data in each batch. Various data augmentation methods are used to overcome possible overfitting when the minority data are oversampled. Two binary classifiers for identifying *OV* and *MEAS* are also trained after balancing the number of images (samples) in each batch.

4.3 CLASSIFICATION RESULTS

The results from one multiclass and three binary classifications are summarized in Table 1. From the testing, classification successfully attains high recall and precision. Overall, the recall in each class is more than 80.0%. Here, we point out two findings regarding the results: First, in binary classifications of *OV* or *MEAS*, precision values are relatively low compared to recall values. This outcome is because the number of testing samples is highly imbalanced, indicating that the number of negative samples is much greater than the number in the target class (*OV* or *MEAS*). In the real-world image set, images in these classes are relatively small portions of the entire sets of *BIN* or *BOUT*. When these classifiers are applied to *OV* and *MEAS*, a large portion of false-positive samples is generated. However, in actual implementation, high recall far outweighs high precision, because it is more critical to identify all of the images in such minority classes without missing them, rather than reducing the occurrence of false-positives. Second, the precision for binary classifications for *BIN/BOUT* and *MEAS* are relatively low as compared to multiclass precision values. The images in each class for multiclass classification have very distinct boundaries of visual features for classes, and share similar appearances in an identical class, causing high classification accuracy. However, for the binary classifications, the images look similar and have less clear boundaries between classes. For example, in some images in the *MEAS* class, the only visual difference is the inclusion of a measurement tool such as a measuring tape or crack gauge. This could be overcome by adding more training images for *MEAS* to generate more features from the difference in the network.

Table 1. Classification results for various image classes

Classification	Multiclass				Binary		Binary	Binary
	<i>BBC</i>	<i>GPS</i>	<i>WATCH</i>	<i>DOC</i>	<i>BIN</i>	<i>BOUT</i>	<i>OV</i>	<i>MEAS</i>
Category								
# of labelled images	16,747	835	320	3,283	6,407	9,650	1,531	690
# of testing images	4,175	224	87	809	1,590	2,423	365	176
Precision	89.7%	78.2%	75.5%	89.7%	86.0%	88.2%	48.9%	50.4%
Recall	97.3%	99.1%	95.4%	97.3%	81.4%	91.3%	91.0%	80.1%

4.4 SAMPLE REPORTS

The ultimate goal of our approach is to incorporate information extracted from a set of images into a single report using trained classifiers. As mentioned in Section 4.1, two sets of images collected after different earthquakes are used for the demonstration of report generation. The images used in the sample report generation are not included in the training process in the previous section. Detailed information for these two collections is presented in Table 2 and the images collected from each building are shown in Figure 4. All images in Figure 4 are presented as square thumbnails for this depiction. Note that each set of images does not necessarily include all classes in the schema. For example, if a digital copy of the drawings is available, the team does not need to collect them as images. The actual appearance of the building can also be found in Google Street Views using GPS information. In Table 2, the definition of damage levels for structural and masonry wall damage used here can be found at datacenterhub.org.

Table 2. Description of the data sets used for generating sample report

Id	Event	Date	Latitude	Longitude	# of images	Structural damage	Masonry wall damage
Set 1	Ecuador earthquake	Jul 16, 2016	-0.594472222	-80.42313889	93	Moderate	Severe
Set 2	Taiwan earthquake	Mar 8, 2016	23.12341667	120.4699444	263	Severe	Moderate

Figure 5 shows the two reports automatically generated from each dataset. All thumbnail images are linked to the original resolution images. Figure 5 shows the images from the labeled datasets described in Section 4.1, although they are not used for training classifiers. Overall, the performance of the classifiers works reasonably well. With such reports, engineers can readily view the general appearance and information collected in the field, as well as being able to reach images-of-special-interest, without the distractions of unrelated images.



(a)



(b)

Figure 4. Two sets of the original images collected from the actual earthquake reconnaissance missions: (a) Ecuador, 2016 and (b) Taiwan, 2016. Note that these images are ordered according to the time the images were taken. All images are rescaled as square thumbnails for the arrangement.

Documentation of Post-Earthquake Reconnaissance Images

Table of Contents

- Summary
- Building and Building Components
 - Outside
 - Inside
- Metadata
 - GPS
 - Drawing
 - Watch
- All Images

Summary



- Date: 2016.07.26
- Time: 13:23 ~ 13:54 (17 minutes)
- The number of images: 19
- GPS/Drawing: 100/0

Building and Building Component

Outside



Inside



Metadata

GPS



Drawing



(a)

Documentation of Post-Earthquake Reconnaissance Images

Table of Contents

- Summary
- Building and Building Components
 - Outside
 - Inside
- Metadata
 - GPS
 - Drawing
 - Watch
- All Images

Summary



- Date: 2016.03.08
- Time: 00:00 ~ 02:02 (122 minutes)
- The number of images: 263
- GPS/Drawing: 100/0

Building and Building Component

Outside



Inside



Metadata

GPS



Drawing



(b)

Figure 5. Sample reports (a) and (b) automatically generated from original images in Figures 4 (a) and (b). MEAS images are marked as red boxes.

5. SUMMARY AND CONCLUSION

In this study, we develop a novel approach based on deep learning for rapidly and autonomously classifying and organizing post-event reconnaissance building images to enable engineers to readily find images-of-special-interest. In the field, this tool is intended to support the field team to augment the data collection effort with proper descriptive information and documentation. We designed an appropriate image schema to provide categories of information that are obtainable from the images, while also being useful for structuring and rapidly browsing through the many images. Using state-of-art image classification algorithms called convolution neural networks, robust classifiers are trained to classify and analyze visual contents (or category) in each image. The information from visual contents and metadata in the image set for each building is then incorporated into a report. In this study, we focus on the developing the schema and demonstrating the approach through the organization of images from buildings during earthquake reconnaissance missions. Using a large volume of real-world images from past missions, we demonstrate the capabilities of our classifiers for image classification. With classifiers successfully trained, two sample reports are generated from two sets of building images. We expect that this strategy will be a breakthrough in browsing and understanding reconnaissance images in a rapid and useful way in the field.

ACKNOWLEDGEMENT

We wish to acknowledge partial support from National Science Foundation under Grant No. NSF 1608762, and the valuable image contributions from CrEEDD (Center for Earthquake Engineering and Disaster Data) at Purdue University (datacenterhub.org), the EUCentre (Pavia, Italy), the Instituto de Ingenieria of UNAM (Mexico), FEMA, and the EERI collections. We also acknowledge the NVIDIA Corporation for the donation of a high-end GPU board.

REFERENCES

- Krizhevsky, A., Sutskever, I., Hinton, G.E. [2012] “Imagenet classification with deep convolutional neural networks,” *Processings of Advances in Neural Information Processing Systems* 25.
- Purdue University, NCREE [2016] “Performance of Reinforced Concrete Buildings in the 2016 Taiwan (Meinong) Earthquake,” <https://datacenterhub.org/resources/14098>.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S. [2015] “ImageNet large scale visual recognition challenge,” *International Journal of Comput Vision*, Vol. 115, pp. 211–52.
- Shah, P., Pujol, S., Puranam, A., Laughery, L. [2015] “Database on Performance of Low-Rise Reinforced Concrete Buildings in the 2015 Nepal Earthquake,” <https://datacenterhub.org/resources/238>.
- Sim, C., Song, C., Skok, N., Irfanoglu, A., Pujol, S., Sozen, M. [2015] “Database of low-rise reinforced concrete buildings with earthquake damage,” <https://datacenterhub.org/resources/123>.
- Sim, C., Villalobos, E., Smith, J.P., Rojas, P., Pujol, S., Puranam, A.Y., Laughery, L., [2016] “Performance of Low-rise Reinforced Concrete Buildings in the 2016 Ecuador Earthquake,” <https://datacenterhub.org/resources/14160>.
- Vedaldi, A., Lenc, K. [2014] “MatConvNet - Convolutional neural networks for MATLAB,” arXiv:1412.4564.

Yeum, C.M., Dyke, S.J., Ramirez, J. [2016a] “Visual Data Classification in Post-Event Building Reconnaissance,” submitted to *Engineering Structures*.

Yeum, C.M., Dyke, S.J., Ramirez, J., Benes, B. [2016b] “Big Visual Data Analytics for Damage Classification in Civil Engineering,” *Transforming the Future of Infrastructure through Smarter Information: Proceedings of the International Conference on Smart Infrastructure and Construction*, 27–29 June 2016, 569-574.

Yeum, C.M. [2016] “Computer Vision-Based Structural Assessment Exploiting Large Volumes of Images,” (Doctoral dissertation), Lyles School of Civil Engineering, Purdue University, West Lafayette, Indiana, USA.