LEARNING FROM EARTHQUAKES USING THE AUTOMATIC RECONNAISSANCE IMAGE ORGANIZER (ARIO)

S. J. Dyke[(1)], X. Liu[(2)], J. Choi[(3)], C. M. Yeum[(4)], J. A. Park[(5)], M. Midwinter[(6)], T. Hacker[(7)], Z. Chu[(8)], J. Ramirez[(9)], R. Poston[(10)], M. Gaillard[(11)], B. Benes[(12)], A. Lenjani[(13)], and X. Zhang[(14)]

[(1)] *Professor, School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA, sdyke@purdue.edu*
[(2)] *PhD Student, School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA, liu1787@purdue.edu*
[(3)] *PhD Student, School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA, choi343@purdue.edu*
[(4)] *Assistant Professor, Department of Civil and Environmental Engineering, University of Waterloo, ON, N2L 3G1, Canada, cmyeum@uwaterloo.ca*
[(5)] *Student, Department of Civil and Environmental Engineering, University of Waterloo, ON, N2L 3G1, Canada, j246park@uwaterloo.ca*
[(6)] *Student, Student, Department of Civil and Environmental Engineering, University of Waterloo, ON, N2L 3G1, Canada, max.midwinter@uwaterloo.ca*
[(7)] *Professor/Interim Department Head, Department Computer and Information Technology, Purdue Polytechnic Institute, Purdue University, West Lafayette, IN 47907, USA, hacker@purdue.edu*
[(8)] *PhD student, Department Computer and Information Technology, Purdue Polytechnic Institute, Purdue University, West Lafayette, IN 47907, USA, chi32@purdue.edu*
[(9)] *Karl H. Kettelhut Professor, Lyles School of Civil Engineering, Purdue University, West Lafayette, IN 47907, USA, ramirez@purdue.edu*
[(10)] *Neil Armstrong Distinguished Visiting Fellow, Purdue University, West Lafayette, IN 47907, USA and Senior Principal, Pivot Engineers, Austin, TX 78746, USA, poston@pivotengineers.com*
[(11)] *PhD student, Department of Computer Graphics Technology, Purdue Polytechnic Institute, Purdue University, West Lafayette, IN 47907, USA, mgaillar@purdue.edu*
[(12)] *Goerge McNelly Professor of Technology and Professor of Computer Science, Purdue University, West Lafayette, IN 47907, USA, bbenes@purdue.edu*
[(13)] *PhD Student, School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA, alenjani@purdue.edu*
[(14)] *PhD student, Lyles School of Civil Engineering, Purdue University, West Lafayette, IN 47907, USA, zhan3794@purdue.edu*

Abstract

In the aftermath of earthquake events, reconnaissance teams are deployed to gather vast amounts of images, moving quickly to capture perishable data to document the performance of infrastructure before they are destroyed. Learning from such data enables engineers to gain new knowledge about the real-world performance of structures. This new knowledge, extracted from such visual data, is critical to mitigate the risks (e.g., damage and loss of life) associated with our built environment in future events. Currently, this learning process is entirely manual, requiring considerable time and expense. Thus, unfortunately, only a tiny portion of these images are shared, curated, and actually utilized. The power of computers and artificial intelligence enables a new approach to organize and catalog such visual data with minimal manual effort. Here we discuss the development and deployment of an organizational system to automate the analysis of large volumes of post-disaster visual data, images. Our application, named the Automated Reconnaissance Image Organizer (ARIO), allows a field engineer to rapidly and automatically categorize their reconnaissance images. ARIO exploits deep convolutional neural networks and trained classifiers, and yields a structured report combined with useful metadata. Classifiers are trained using our ground-truth visual database that includes over 140,000 images from past earthquake reconnaissance missions to study post-disaster buildings in the field. Here we discuss the novel deployment of the ARIO application within a cloud-based system that we named VISER (Visual Structural Expertise Replicator), a comprehensive cloud-based visual data analytics system with a novel Netflix-inspired technical search capability. Field engineers can exploit this research and our application to search an image repository for visual content. We anticipate that these tools will empower engineers to more rapidly learn new lessons from earthquakes using reconnaissance data.

Keywords: post-event reconnaissance; big data; deep convolution neural networks; data analytics.

1. Introduction

Perishable data that are collected during post-event reconnaissance missions has historically been quite informative. Learning from such data enables engineers to gain new knowledge about the real-world performance of structures, and such investigations have regularly resulted in observations of behaviors not possible under the constraints of laboratory conditions. Thus, this new knowledge, extracted from visual data, is critical to mitigate the risks (e.g., damage and loss of life) associated with our built environment in future events. These visual data serve as evidence to motivate new research directions, and ultimately may lead to code changes. Damaged buildings and components provide learning tools; lessons to be learned from the buildings that do not experience damage are just as important. Gathering such evidence is so important that dozens of teams of engineers, at great expense, typically descend on a site after an event to acquire these data, and the National Sciencce Foundation has funded many projects to coordinate collection of these data.

The predominant method to collect this important field data is through photographs. While the cameras being used have become much more powerful and more prevalent than in the past, and modern cameras are easier to carry and manipulate, the underlying approach to capturing this data is essentially the same as it was several decades ago [1, 2]. Although newer sensor platforms are available, such as LiDAR and ultrasonic, taking photographs is simply the most natural way to collect such visual evidence [3–6]. However, only a small portion of the images collected are actually being used for research and extracting lessons both onsite as well as in longer-term studies. Even less data are being used for making decisions in the field regarding optimizing the collection of perishable data.

Currently, this learning process is done entirely manually, requiring considerable time and effort. Although the images are numerous, existing data repositories today are still quite limited in their ability to make these valuable and voluminous data discoverable and searchable. Currently, documentation and curation of the images takes a great deal of time, and other researchers cannot easily use the data even when they are shared. Thus, the benefits derived from the data collection process are essentially limited to the data collector. Unfortunately, only a tiny portion of these images are shared, curated, and actually utilized. Even when they are shared in open repositories, difficulties still arise in the use of the images collected by another researcher because search is difficult and much of the context is missing.

Computer vision and machine learning have the potential to transform the way that data collection and structural engineering research based on such visual data is conducted. ARIO, the *Automated Reconnaissance Image Organizer*, has been developed by the authors with this goal in mind. We have made significant progress toward robust use of these powerful methods for various purposes, *e.g*, to recognize real-world scenes, to detect structural damage, and to classify and annotate each image. ARIO enables the field engineer to rapidly and automatically categorize their post-event reconnaissance images. Using more than 140,000 images collected from reinforced-concrete buildings during past earthquake reconnaissance missions, we have established a labeled, ground-truth visual database. The authors used the database to train and validate robust classifiers across a range of categories that are useful within the domain of structural engineering [7–11]. By exploiting these carefully trained, domain-oriented classifiers, ARIO yields a structured report combined with useful metadata, and integrates the capability to filter the data by class to look across time and region [12, 13].

This research and development drives the field of disaster informatics into new territory by facilitating automated sifting, classification and analysis of *hundreds of thousands of images of real disaster scenes*. The engineering researcher is able to readily extract important images and information from large volumes of field data. Moreover, it sets in motion a new paradigm for investigation across the multidisciplinary field of resilience and disaster research by enabling visual data obtained after real events, and gathered over decades by thousands of individuals around the world, to be analyzed, filtered, described, and repurposed for generating new knowledge and making our communities more resilient and sustainable.

2. Description of the ARIO Application

ARIO is intended to rapidly and automatically categorize large volumes of post-event building reconnaissance images. The engineer can use it in the field to assess the information provided from the data collected during a period in the field, or can use it in the office to automatically prepare reports describing the data collected to share with others. It should be accessible for field engineers to upload sets of building data over the internet, with remote analysis and classification to efficiently organize the data into a report. The application is illustrated in Fig. 1.

ARIO is made possible through the integration of three major recent achievements: (i) automated image classification based on visual contents [11], (ii) domain-oriented taxonomy [8, 13], and (iii) cloud based infrastructure. First, the ability to train robust scene classifiers has been developed by the authors for the types of unstructured scenes collected in reconnaissance missions to support categorization of images based on visual contents. Such classifiers are trained and validated with our ground-truth database containing tens of thousands of images of reinforced concrete buildings captured from past reconnaissance missions. Second, a suitable domain-oriented taxonomy has been created to capture the categories that are most useful for structural engineers. This development is based on the experience of field engineers, the types of images that they actually collect, and the reasons engineers would aim to sort and filter images to extract new knowledge. The taxonomy is employed to manually annotate images from past events, yielding a sufficiently large training dataset. Last, a cloud-based visual analytics platform is established to integrate these capabilities in a tool that can be accessed through a web browser, and thus provide the means for multiple users to eventually be able to share these data and reports as a web service.

Fig. 1 – Overview of ARIO: Rapid, automated reconnaissance image classification, data sets uploaded from the field or the office over the web; report generation; and sorting and filtering.

2.1 Robust reconnaissance image classification

Automated scene classification enables the extraction of visual contents from images based on a pre-defined taxonomy. We first generated a ground-truth building reconnaissance database which contains over 140,000 images gathered from publicly available reconnaissance image repositories [14–17]. Both the training and validation of the classfiers requires these labeled data to fine-tune robust classifiers and perform the fine-tuning of parameters. This research leveraged the vast amount of work done to etablish high quality pre-trained convolutional neural networks (CNNs) [18, 19]. However, once the robust classifiers are trained on past data, we have found them to be quite robust for application to a broad range of data sources as documented and validated in [8, 20–21]. Since scenes are understood by low-dimensional features (e.g. general shape, colors, or compositions), so classification can be performed rapidly in about 1 second per image. Thus, efficient analysis of large-scale visual data is quite feasible on a cloud-based platform.

To support classification of post-earthquake reconnaissance scenes, we have been successful in designing classifiers for several image categories such as building components (e.g., wall, columns, or beams), damage level (e.g., cracking, spalling, or collapse), and image location (e.g. building interior or exterior). Images that contain useful metadata are also commonly collected during a reconnaissance mission (e.g., images of GPS, drawings, watches, or measurements) and these can be classified and annotated for easy access.

2.2 Taxonomy for the categories in structural engineering

Field engineers were consulted to develop appropriate visual data categories along with an associated hierarchical structure to support the domain needs of the application. The categorical classes selected for the taxonomy must be appropriate and useful for the specific application to be supported. However, classification can only differentiate between classes if images in each classes are visually distinguishable. A class hierarchy is designed here to enable both efficiency and accuracy in performing classification. Thus, the relationship between the classes we have defined may not be mutually exclusive, meaning that a given image may actually be categorized within multiple classes. Note that the multilabel classification algorithm we have implemented does allow multiple probable classes without reducing classification accuracy [22, 23].

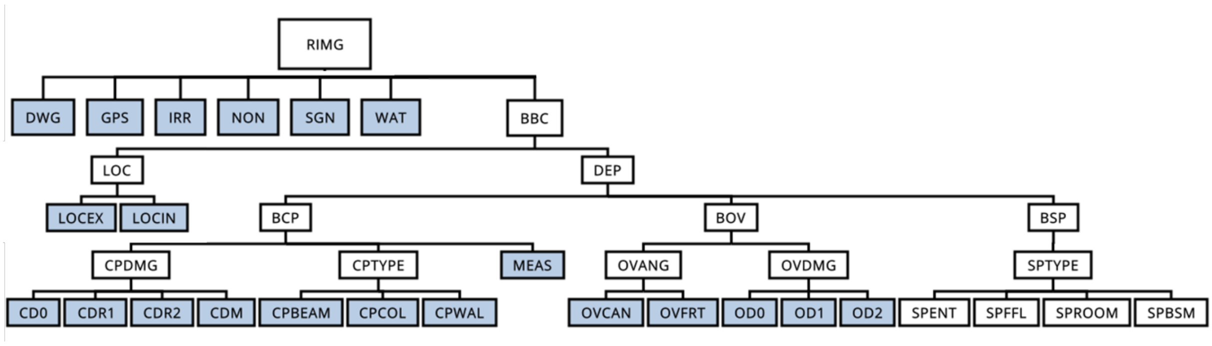

Fig 2 – Hierarchy of the classes included in the ARIO schema

To successfully classify the images based on our taxonomy, a clear definition is necessary for each category. Definitions are essential, as they provide guidance to experts, human annotators, who provide consistent ground truth data that are suitable for training meaningful classes. Moreover, providing clear boundaries that can be used to distinguish the visual features of images in the different classes will help to improve the robustness and accuracy of the classification results. We have defined each category carefully, and these definitions are then used by human annotators to label images to form our ground truth database. ARIO classifiers are built based on both this taxonomy and the ground truth database.

The taxonomy generated for reconnaissance images (RIMG) in this application is shown in Fig. 2. Boxes shown in solid blue (filled with color) are the subset of classes actually used for directly categorizing each image. A trained classifier must be created for each of these classes. The remaining classes shown with white boxes are then included to complete the taxonomy, but are only indirectly classified.

As mentioned earlier, some images collected in the field contain important metadata, such as the categories we define as: Drawing (DWG), GPS, non-building component (NON), sign (SGN), watch (WAT), and measurement (MEAS). Most of the images though fall into the category building and building components (BBC), and then dig into the purpose and type of image collected. For instance, the location of the image is captured in classes for image Location (LOC), building exterior (LOCEX), building interior (LOCIN). The visual scene contained in the image could be either a building overview (BOV) showing the entire building from the outside, or a building component (BCP) showing some portion of a building from either the inside or outside. These are all contained within the indirect category we call depth (DEP).

BOV images can be classified by the orientation of the camera, as in canonical view (OVCAN), overview image captured angle (OVANG), and front view (OVFRT). They can also be classified within overview damage (OVDMG), according to the level of damage present, as minor overall damage (OD1), moderate overall damage (OD2), and severe overall damage (OD3).

BCP images can be classified as column (CPCOL), beam (CPBEAM), or wall (CPWALL) which define the component type (CPTYPE). Furthermore, they can be organized in terms of their component damage (CPDMG), for instance as minor/no damage (CD), moderate concrete damage (CDR1), severe concrete damage (CDR2), or masonry damage (CDM).

There are also irrelevant (IRR) images, that are not specifically intended for building performance assesssment. These are frequently collected in the field, but are best filtered out and put aside for possible use to support other purposes.

This taxonomy is designed with the engineering user in mind.

2.3 Cloud-based visual analytics platform establishment

To create a cloud-based visual analytics platform for the civil engineering community, our first step was to migrate the ARIO tool to run within docker containers. We also created a docker repository to store and distribute ARIO containers for docker installations. As an online web-based tool, ARIO's deployment provides a domain science specific cloud-based data analytics service that seeks to integrate *application-as-a-service* and *data-as-a-service*. In terms of application-as-a-service, ARIO reduces the effort required to manually label post-disaster visual data. Users can build reports that are based on uploaded post-disaster images into the ARIO database that are classified as a step that takes place in parallel with uploading the data set from a single building. Also, users can download an existing report into a zip file to view it offline. As for image classification, ARIO automatically classifies the uploaded post-disaster images and labels the images with its category or categories if an image is classified as multiple categories. Moreover, users can use Google Map functionality to show the location of the event. Additionally, as for data-as-a-service, we set up a FreeNAS storage server to store visual data of the ARIO. The FreeNAS also serves as the start point of the data backbone.

Fig. 3 shows the detailed deployment of ARIO on our infrastructure. We put the ARIO system behind a pfSense firewall and applied a load-balancer to distribute the requests to several servers. We utilized Docker to achieve the expandability and scaled the system up to multiple servers. By migrating the ARIO application as well as the system's Microsoft SQL Server database to running in Docker containers, the application can be distributed on new machines quickly and easily. Furthermore, we set up a separate server machine that handles classification tasks purely. The requests coming to the classification server are also distributed to multiple Docker containers that run the neural network model. The distribution of requests is completed with NGINX.

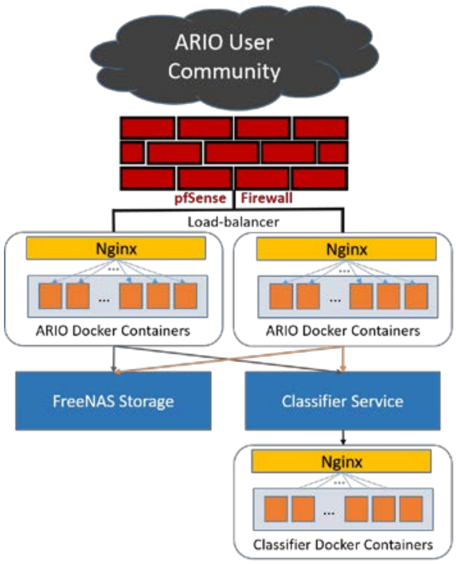

Fig. 3 – ARIO requires a secure, distributed, shared cyberinfrastructure that can be used to process a large volume of data and provide access to multiple users at a time.

3. ARIO Deployment

ARIO is designed to execute through a web browser on local computers that may have various operating systems including Windows, Mac OS or Linux. It is also accessible on mobile devices, such as iOS or Android in smartphones. The web-based ARIO deployment allows the processing to run at a fast speed and behind the scenes, and it is fully automated.

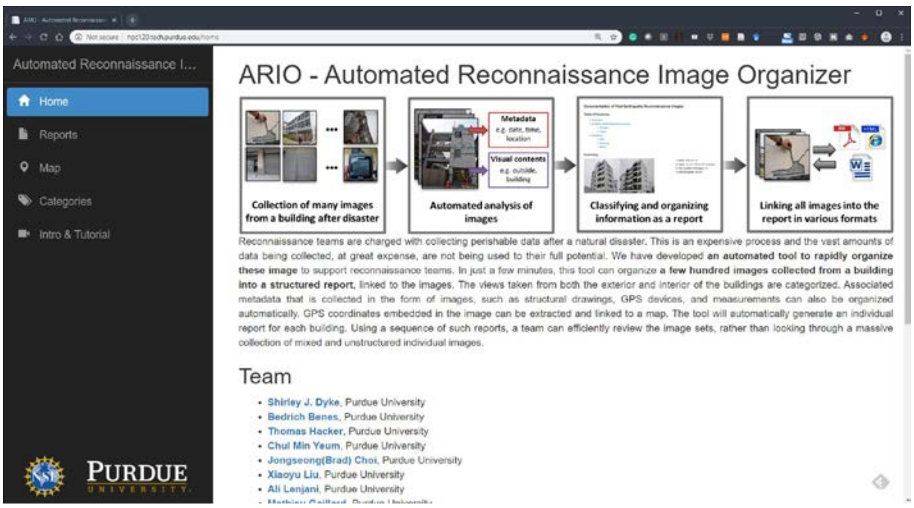

Fig. 4 – Main page of the ARIO application in a web browser

The main landing page of ARIO is shown in Fig. 4. Some general information is provided on this page including an overview of the application and the workflow to use the application, as well as the many data sources that enabled this development. Five main functions are shown on the left sidebar to enable easy navigation to use the functionalities of ARIO. The ‘Home’ menu item provides a link to allow upload of images to the server to generate a new report, as well as to browse past reports that were previously generated. The ‘Map’ menu item displays all of the locations where data collection took place and a report was generated. This simple capabilty is performed by automatically extracting GPS information from the uploaded images. The ‘Categories’ selection defines the list of categories available for filtering images. By choosing each category, the images are filtered and highlighted based on the classified reports; those not in a selected category are dimmed. The user may use these to look across reports and compare images within a given class. Finally, 'Intro & Tutorial' offers a user manual and an intuitive tutorial about how to begin using ARIO.

The main function ARIO is to provide the ability to automatically categorize reconnaissance images submitted by user and generate a report. The time required for this process will depend on the volume of images being processed, but generally requires about a second of processing time per image. Once the processing is complete, the user will immediately see the resulting report on the current browser window. ARIO is designed to group images collected in one building within one report. If user intends to categorize images taken from multiple buildings, that user can simply create several new reports and upload the sets of images separately.

Next we run through the various functionalities and show examples from previously generated building reconnaissance reports and filtering operations to provide the reader with an idea of how the ARIO application can be used.

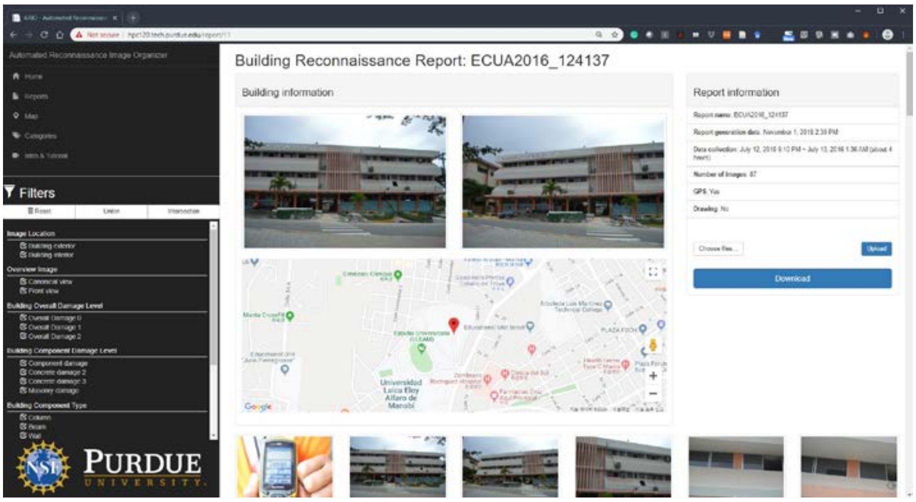

Fig. 5 – An example view of a report generated, showing a couple building overview images, the location information on a Google map, and automatically extracted information from the photo metadata.

An example of a generated report is shown in Fig. 5. At the top, two overview images are selected and shown to help user easily recognize the building. ARIO also provides additional information about these images besides the visual contents. Under the two overview images, the location of this building is marked as a pinpoint on Google map. Thus, a user can quickly retrieve images from sometime before the earthquake using this Google map or by navigating to a Street view. General information about both the building and the report are also provided near the top of the page, including the data these images were collected, the number of the images in this report, and when possible, if GPS and structural drawings are available. All of this information is automatically extracted from the image set and provided to the user.

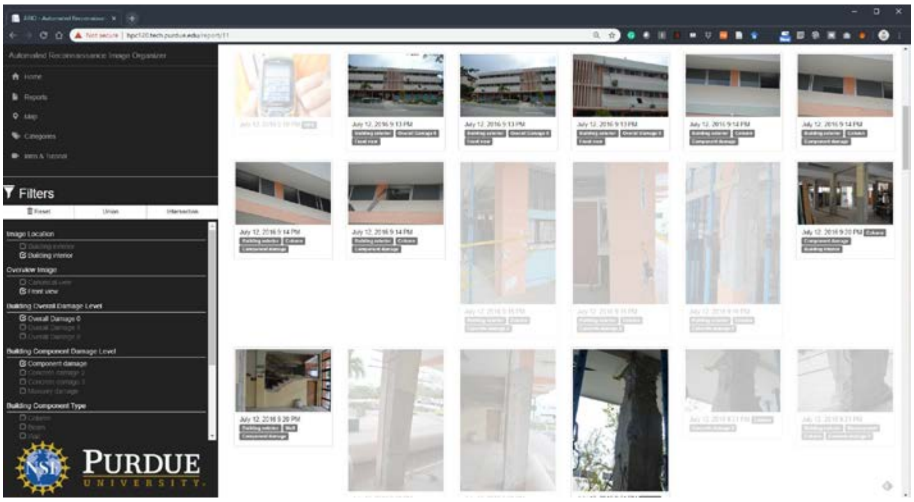

Fig. 6 – Image filtering function for categorizing images enabled by our pre-trained classifier

All of the uploaded images are listed below the metadata information, as captured in Fig. 6. Under each image, the classified results are shown as greyed labels. Our classification schema is designed to assign multiple category labels to each image, so the number of labels under each image may be different. To help a user easily browse through these classification results, ARIO integrates an interactive tool to filter the desired categories of images. To use the filter in the left column of the window, the user selects the categories of greatest interest. More than one category can be selected. Then, the images which are not in the selected categories are shown as dimmed, only the images with selected categories are highlighted. The filter also supports the intersection or union of multiple selected categories. In this way, the user is able to extract important images or information among the large volume of data in this report in just a few seconds. For example, a user could easily identify images that show severe column damage without manual searching through a large image set. Finally, ARIO allows the user to download the report to the user's local storage.

4. Future Expansion: VISER

(a) (b)

Fig. 7 – Future applications to produce VISER: (a) overall concept for Netflix-like search capability, an image is used to initiate the search, various databases are subscribed to, and similar images are provided; (b) automated reconstruction of path mapping for indoor reconnaissance missions, localizing images and linking them to locations on a structural drawing

To extend the availability of ARIO, our plan for future research is to develop and integrate useful functionalities. Our first vision is to construct a science-oriented, visual data service that has a science-driven, Netflix-like search capability to interrogate data sets based on visual content, as illustrated in Fig 7 (a). These capabilities will enable data and applications to be brought to the user seamlessly. The user can choose the level of engagement by selecting the data repositories to connect to and the applications that he/she would like to access, and processing will take place using a virtual cloud, spun up on-demand, and designed by each user for a particular search and extraction of knowledge. Human experts, engineers and scientists, would then be able to evolve and grow this system, automatically building within it the capabilities to retrieve more and more meaningful visual contents correlated to domain knowledge, and publish those search results and new classifiers to share with the community.

Our second vision is to have the ability to reconstruct the path of an inspector through a site with only image data. Structural drawings are essential as a record of the structural information needed to understand the consequence of events and extract valuable lessons to improve future performance. The data collector can carry one more camera during the mission. As shown in Fig 7. (b), our technique would automatically rebuild the path of the inspector through the indoor environment using the image sequence captured by the additional required camera. And we will use our previously developed drawing reconstruction tool to rebuild the overall version of the drawing based on partial drawing images. Later, the path will be overlaid on the reconstructed drawing as the final output of the technique.

5. Conclusions

A great many opportunities to extract new knowledge and learn more about the performance of our buildings are being lost. In the aftermath of earthquake events, reconnaissance teams are deployed to gather vast amounts of images. Despite the great deal of time and effort spent to capture these perishable data before they are destroyed, only a small portion of these images are really being utilized to learn from earthquakes. Machine learning, when used carefully, can greatly assist with the analysis, organization, filtering and sorting of these data. Here we harness the power of computers and artificial intelligence to organize and catalog such visual data with minimal manual effort. Our application, called the Automated Reconnaissance Image Organizer (ARIO), allows a field engineer to rapidly and automatically categorize large volumes of post-disaster visual data, images and generate a report.

At this time ARIO is intended for earthquake reconnaissance missions due to the large number of images typically collected after an event. However, it could similarly be used for categories more apporpiate for hurricane, fire or flood damage with the proper ground truth databased and well-designed taxonomy.

Future efforts are focused on integrating this application into VISER, a cloud-based system that would allow various data sources to be combined to utilize these classification and filtering functionalities to look across data sources, as well as other applications under development.

5. Acknowledgements

We would like to acknowledge support from the US National Science Foundation under Grant No. NSF 1835473, and our data providers including the CrEEDD (Center for Earthquake Engineering and Disaster Data) at Purdue University; the EU Centre in Pavia, Italy; The Instituto de Ingenieria of UNAM in Mexico City; the EERI Clearinghouse, USA.

6. References

[1] Gur T, Pay AC, Ramirez JA, Sozen MA, Johnson AM, Irfanoglu A, Bobet A (2009): Performance of school buildings in Turkey during the 1999 Düzce and the 2003 Bingöl Earthquakes, *Earthquake Spectra,* **25** (2), 239-256

[2] Purdue University and NCREE (National Center for Research on Earthquake Engineering) (2016). Performance of reinforced concrete buildings in the 2016 Taiwan (Meinong) earthquake. Accessed October 28, 2017. https://datacenterhub.org/resources/14098.

[3] Liao Y, Wood RL, Mohammadi ME, Devkota K, Wittich CE (2019): Damage Assessment of a Sixteen Story Building Following the 2017 Central Mexico Earthquake. *Civil Engineering Faculty Publications.* **157**.

[4] Adriano B, Koshimura S, Karimzadeh S, Matsuoka M, Koch M (2018): Damage Mapping After the 2017 Puebla Earthquake in Mexico Using High-Resolution Alos2 Palsar2 Data. *IEEE International Geoscience and Remote Sensing Symposium, 22–27 July 2018*, Valencia, Spain.

[5] Villalobos E, Sim C, Smith-Pardo JP, Rojas P, Pujol S, Kreger ME (2018): The 16 April 2016 Ecuador Earthquake Damage Assessment Survey. *Earthquake Spectra*, **34** (3), 1201-1217.

[6] Shah P, Pujol S, Puranam A, Laughery L (2015): Accessed October 9, 2017. Database on performance of low-rise reinforced concrete buildings in the 2015 Nepal earthquake.

[7] Yeum CM, Dyke SJ, Ramirez JA, Benes B (2016): Big visual data analytics for damage classification in civil engineering. *International Conference on Smart Infrastructure and Construction, 27–29 June 2016*, London, UK.

[8] Yeum CM., Dyke SJ, Ramirez JA, Hacker T, Pujol S, Sim C (2017): Annotation of image data from disaster reconnaissance. *16th World Conference on Earthquake Engineering, 9–13 January 2017*, Santiago, Chile.

[9] Yeum CM, Dyke SJ, Benes B, Hacker T, Ramirez JA, Lund A, Pujol S (2017): Rapid, automated image classification for documentation. *7th Conference on Advances in Experimental Structural Engineering, 6–8 September 2017,* Pavia, Italy.

[10] Yeum CM, Dyke SJ, Benes B, Hacker T, Ramirez JA, Lund A, Sim S, Pujol S (2018): Automating damage classification in post-earthquake building images. *11th National Conference on Earthquake Engineering, 25–29 June 2018*, Los Angeles, CA, USA.

[11] Yeum CM, Dyke SJ, Ramirez JA (2018): Visual data classification in post-event building reconnaissance. *Engineering Structures*, **155**, 16-24.

[12] Yeum CM, Lund A, Dyke SJ, Ramirez JA (2019): Automated recovery of structural drawing images collected from postdisaster reconnaissance. *Journal of Computing in Civil Engineering*, **33** (1), 04018056.

[13] Yeum CM, Dyke SJ, Benes B, Hacker T, Ramirez JA, Lund A, Pujol S (2019): Postevent reconnaissance image documentation using automated classification. *Journal of Performance of Constructed Facilities*, **33** (1), 04018103.

[14] Datacenterhub (2014). Accessed January 10, 2019. http://datacenterhub.org.

[15] DesignSafe-CI (2016). Rapid experimental facility. Accessed October 9, 2018. https://rapid.designsafe-ci.org.

[16] EERI (Earthquake Engineering Research Institute) (2009). Earthquake clearinghouse. Accessed January 8, 2020. www.eqclearinghouse.org.

[17] UC CEISMIC (2011). Canterbury earthquake digital archive. Accessed October 9, 2017. www.ceismic.org.mz

[18] Krizhevsky A, Sutskever I, Hinton GE (2012): Imagenet classification with deep convolutional neural networks. *25th International Conference on Advances in Neural Information Processing Systems,* Lake Tahoe, Nevada, USA.

[19] Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, Berg AC, Li F (2015): Imagenet large scale visual recognition challenge. *International journal of computer vision*, **115** (3), 211-252.

[20] Sim C, Song C, Skok N, Irfanoglu A, Pujol S, Sozen, M (2015). Database of low-rise reinforced concrete buildings with earthquake damage. Accessed October 9, 2018. https://datacenterhub.org/resources/123.

[21] Sim C, Villalobos E, Smith JP, Rojas P, Pujol S, Puranam AY, Laughery L (2017). Performance of low-rise reinforced concrete buildings in the 2016 Ecuador earthquake. Accessed October 9, 2018. https://datacenterhub.org/resources/14160

[22] Zhou B, Lapedriza A, Xiao J, Torralba A, Oliva A (2014): Learning deep features for scene recognition using places database. *27th Conference on Advances in Neural Information Processing Systems*, Lake Tahoe, NV, USA.

[23] Vedali A, Lenc K (2015): MatConvNet: Convolutional neural networks for MATLAB. *23rd ACM International Conference on Multimedia*, New York, NY, USA.