



IDENTIFICATION OF DIAPHRAGM FLEXIBILITY OF UNREINFORCED MASONRY BUILDING USING CONVOLUTIONAL NEURAL NETWORKS

A.B. Acevedo⁽¹⁾, D. Rueda-Plata⁽²⁾, D. González⁽³⁾, J.C. Duque⁽⁴⁾, R. Ramos-Pollán⁽⁵⁾

⁽¹⁾ Professor of Civil Engineering, Applied Mechanics Research Group, Universidad EAFIT, aaceved14@eafit.edu.co

⁽²⁾ Master Student of Applied Mathematics, CAGE Research Group, Universidad Industrial de Santander, ing.diegorueda@gmail.com

⁽³⁾ Doctoral Student of Engineering, Applied Mechanics Research Group, Universidad EAFIT, dgonza26@eafit.edu.co

⁽⁴⁾ Research in Spatial Economics (RISE-group), Department of Mathematical Sciences, Universidad EAFIT, jduquec1@eafit.edu.co

⁽⁵⁾ Professor, Universidad de Antioquia, raul.ramos@udea.edu.co

Abstract

The classification of the building stock for risk assessment can become a great challenge as it requires a description of each building in terms of its capacity to resist a given hazard. When evaluating the seismic risk, the assessment requires a typology classification based on its capacity to resist lateral loads. Unreinforced masonry, a building typology in which lateral loads are sustained by masonry walls without any form of reinforcement, is highly vulnerable to earthquakes. While in Europe this type of building is commonly found in historical centers, in developing countries like South American ones, unreinforced masonry may represent a high proportion of the building stock. For example, 74% of the residential stock in the metropolitan area of Medellín (Colombia) is unreinforced masonry. In these cases, it is necessary to go beyond the mere classification as "unreinforced masonry" by including additional characteristics such as the flexibility of the floor/roof diaphragms, which plays an important role in the distribution of the lateral forces to the vertical elements of the lateral load-resisting system. In addition, a rigid diaphragm allows the structure to have synchronized displacements during an earthquake. As the availability of street-level imagery becomes widespread, a great potential emerges to support detailed city-wide risk assessment in a cost-effective manner with image analytics. In this paper, we explore this potential by evaluating five state-of-the-art architectures of convolutional neural networks (CNN) to classify images of one-story unreinforced masonry buildings of the metropolitan area of Medellín according to the flexibility of the roof diaphragm. Over a thousand images were manually gathered from Google Street View and properly labeled by experts to feed the CNN training process. To select the best performing architecture we followed a conservative, risk-averse approach, by minimizing the risk of classifying a flexible diaphragm as rigid on validation images (not seen during the training stage). In our process we (1) split the data in train/validation/test subsets, (2) train each CNN architecture on train data, (3) select the best according to their performance in validation data, and (4) use test data to obtain an unbiased performance estimation of our process as a whole. The best performing model, Vgg19, reaches an accuracy of 80%, a precision of 88%, and a recall of 84% on test data. The accuracy level is considerably higher for flexible diaphragms than for rigid diaphragms, 85% vs. 71%. We also found that there is a direct relationship between the level of certainty assigned by the experts when classifying an image and the probability assigned by the CNN.

Keywords: unreinforced masonry; diaphragm; convolutional neural networks; seismic risk; risk assessment.



1. Introduction

Evaluation of the risk due to natural disasters is fundamental for disaster management as it quantifies the possibility of economic and human affectations in a given region. In order to assess the risk three models are needed: a hazard assessment that characterizes the expected physical event (e.g., earthquake, flood, fire, etc.) in terms of frequency and size; an exposure model that describes the characteristics of the assets (people and infrastructure) in the region under study; and a vulnerability model that relates the physical event intensity to the assets capability of sustain such forces. Each of the three aforementioned aspects is fundamental to the risk assessment as the reliability of the results is highly dependent on the quality of the three input models.

This paper focuses on the exposure model for seismic risk assessment of buildings; i.e., the description of the buildings' characteristics based on their capacity to resist lateral loads. More specifically, the paper concentrates on the classification of one-story unreinforced masonry buildings according to its roof flexibility (rigid or flexible diaphragm). Seismic building capacity is defined by factors such as the material and type of the lateral load resisting system, ductility, number of stories, date of construction, the shape of the building plan, etc. For example, the recent Global Earthquake Risk Model [1] classified buildings based on the material of the lateral load-resisting system (LLRS), the type of LLRS, the number of stories, and the building ductility. On the other hand, Hazus [2] building inventory classification uses the basic structural system and three building height subclasses (low-rise, mid-rise and high-rise). In the aforementioned studies unreinforced masonry buildings constitute a building classification that identifies a non-ductile structure in which the LLRS are masonry walls, without any type of reinforcement. This classification is, in most cases, good enough for the evaluation of the seismic risk. Nonetheless, in regions where this building typology represents more than half of the building stock, it may be desirable to include additional features—such as diaphragm flexibility—in such a way that a better description of the building seismic capacity can be achieved.

We consider the diaphragm flexibility a relevant characteristic as it plays an important role in the distribution of the lateral forces to the vertical elements of the LLRS. Furthermore, a rigid diaphragm allows the structure to have synchronized displacements during an earthquake, which implies a better seismic response. The inclusion of diaphragm flexibility as a parameter on the building exposure model requires the manual classification of unreinforced masonry buildings based on their diaphragm characteristics: flexible or rigid. To decrease the additional amount of work on the exposure model development due to the inclusion of the flexible/rigid diaphragm characteristic, we explore in this paper the potential of use convolutional neural networks (CNN) to classify street-level imagery of one-story unreinforced masonry buildings according to the flexibility (rigid or flexible) of the roof diaphragm. We use the Metropolitan Area of Medellín (henceforth referred to as Medellín) as a case study. Medellín is the second-largest conurbation of Colombia and 74% of its residential building stock is composed of unreinforced masonry buildings [3].

As a first step of the present study, we manually gathered from Google Street View over one thousand images of one-story unreinforced masonry buildings of Medellín. Each image was labelled by experts as flexible or rigid roof diaphragm. Details of the database are given in section 2. The database was used for the evaluation of five state-of-the-art architectures of CNNs. For each architecture data was (1) split in train/validation subsets, (2) the CNN architecture was trained and (3) the performance was measured. We selected the best architecture by a conservative, risk-averse approach that consists of minimizing the risk of classifying a flexible diaphragm as rigid. Section 3 presents the methodology used for the selection of the best CNN architecture. Results of the best CNN architecture are presented in Section 4, in which its performance is measured in terms of statistical parameters. Finally, section 5 presents the main conclusions of this work.

2. Unreinforced masonry buildings of Medellín and its metropolitan area: database

Unreinforced masonry (MUR) is a building type in which walls are made of masonry without any form of reinforcement. Masonry refers to individual units of brick, stone, concrete block, adobe, etc. that are joint together by mortar. Unreinforced masonry buildings are non-ductile structures, i.e., they have a reduced capacity to resist lateral deformation; therefore, they should not be used in medium to high seismic hazard zones. Regardless of their inability to sustain seismic loads, unreinforced masonry buildings represent and



important percentage of the Latin America building stock. According to [4] 31% of the residential buildings of the Andean region of South America are MUR buildings. The percentage significantly increases for Medellín, in which different studies have reported the percentage of MUR buildings between 50% and 70% approximately [3, 5, 6].

When a building is subjected to an earthquake lateral loads are imposed on the structure. Such loads must be transmitted to the lateral load-resisting system (LLRS), which, in the case of a MUR building, is the unreinforced masonry walls. The horizontal system that transmits the lateral forces to the LLRS is known as a diaphragm which can be a roof, floor or any other type of horizontal bracing. There are two types of diaphragms: rigid and flexible, which are defined based on the relative stiffness of the vertical elements versus the horizontal elements. A rigid diaphragm is assumed to not deform itself under lateral loads, it is able to distribute the seismic loads to the LLRS, and it allows the structure to have synchronized displacement during an earthquake. On the other hand, a flexible diaphragm it is not able to distribute torsional and rotational forces to the LLRS. A building with a rigid diaphragm will have a better seismic performance than a building with flexible diaphragm.

A seismic hazard assessment usually does not classify MUR buildings considering the type of floor/roof diaphragm. Nonetheless, when the percentage of MUR buildings is high, as in the study area of Medellín, it may be important to discriminate by the type of diaphragm as the seismic performance of a MUR building is highly dependent on the diaphragm flexibility. Vulnerability functions for MUR buildings of Medellín that include the type of diaphragm have been developed by [7]. A vulnerability function—input data for a seismic risk assessment—relates the expected damage with an intensity measure of the ground motion generated by the earthquake. The functions generated by [7] indicate a smaller vulnerability of the MUR buildings that have rigid diaphragms when compared to the vulnerability of MUR buildings with flexible diaphragm.

In this paper we explore the potential use of convolutional neural networks (CNN) to classify images of one-story MUR buildings according to the flexibility (rigid or flexible) of the roof diaphragm. For such purpose we manually collected more than one thousand street-level images of one-story buildings of Medellín and La Estrella—a municipality of the metropolitan area of Medellín—with the use of the Google Street View application. Each image was labeled by experts based on its roof diaphragm as flexible or rigid. In addition, a confidence level was assigned to each classification, i.e., it was recorded the certainty level (100%, 75% and 50% certainty) of the expert when defining the type of roof diaphragm of each building. The uncertainty in the assignation of the type of roof diaphragm arises from the fact that the classification is done based on an image and, therefore, there is not access to the interior of the building. Regardless such disadvantage, remote building surveys have been used for the development of seismic exposure models [3, 8] and it has become a useful tool that reduces the amount of resources involved in the development of such models.

The database includes a total of 1,122 images distributed as 745 of buildings with flexible diaphragm (66% of the total) and 377 buildings with rigid diaphragm (34% of the total). The unbalanced number of buildings of each category reflects the building type distribution in the city: it is usual that a one-story MUR building with a rigid diaphragm will evolve in a building with two or more number of stories; therefore, it is more common to find one-story buildings with flexible diaphragm. Fig. 1 presents the distribution of the surveyed buildings of the database as well as examples of buildings with flexible and rigid diaphragm.

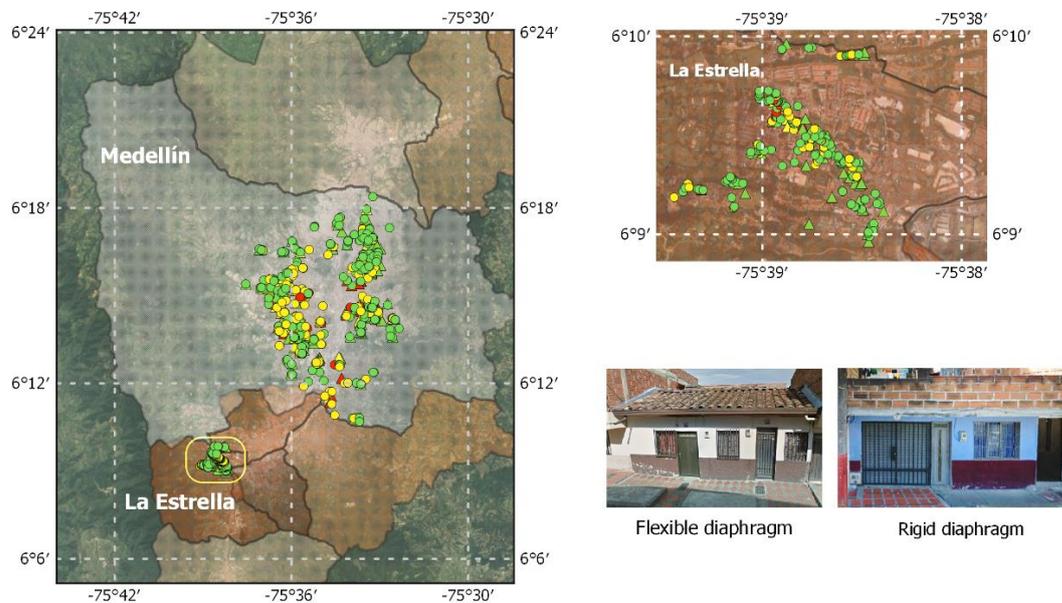


Fig. 1 – Distribution of survived buildings and examples of roof diaphragm type. Triangles: rigid diaphragm; circles: flexible diaphragm. Label color indicates expert certainty level (green: 100% certainty; yellow: 75% certainty; red: 50% certainty)

3. Deep learning

In the last decade, Deep Learning methods have considerably improved the state-of-the-art in several perceptual challenges, particularly for computer vision problems [9]. These models learn specific image features directly from training data without the need for hand-engineered features. The CNN is a specialized deep learning algorithm that exploits spatial or time correlation in the data [10]. In a CNN, images pass through consecutive layers arranged in a hierarchical system, as the image moves forward through the network, the most optimal features are learned and used to classify an image, recognize objects, etc. The architecture of a CNN (Fig. 2) consists of multiple combinations of convolutional layers and pooling layers followed by supporting or regulatory units (such as dropout, batch-normalization, etc.).

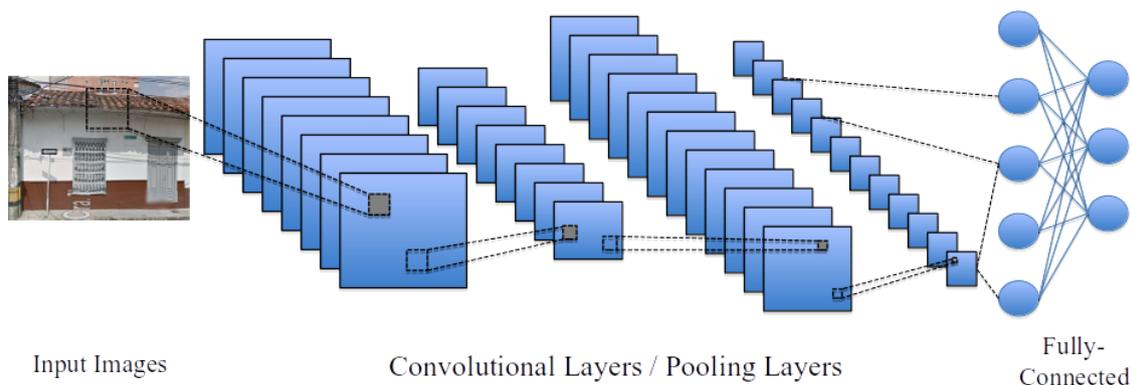


Fig. 2 – Example of a CNN Architecture



The Convolutional layer is composed of a set of kernels, each one convolves on the images using a small receptive field and are activated when a specific feature pattern has been found in the image [10]. Pooling layers extract a combination of features which are invariant to translations and small distortions, while reducing the number of parameters and complexity in the network. Finally, the network is completed with fully-connected layers (FC), these learn non-linear combinations of the high-level features as represented by the output of the convolutional layer, which are then used for the classification of data. CNNs learn hierarchies of visual features (Fig. 3), from simple visual primitives (edges, gradients, borders) in earlier layers of the network, to complex high-level features in later layers. These richer features are composed using combinations of basic structures from previous layers as building blocks.

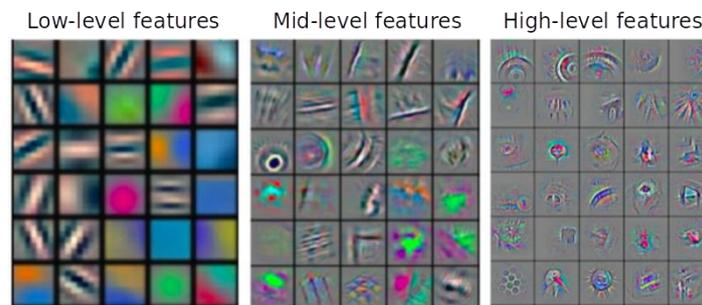


Fig. 3 – Example of hierarchical features

Other than simulation-based approaches (see [11, 12]) and classical statistical or modeling methods [13, 14, 15], there has been little use of machine learning methods for seismic risk assessment. Certain machine learning approaches have been based on applying standard algorithms to data acquired from sensors or databases (public registries, etc.) such as the works in [16, 17, 18]. However, various image processing approaches have been proposed in related areas, including earthquake damage estimation [19], exposure estimation [20, 21], landslide risk assessment [22], etc. Nevertheless, the vast majority of these approaches are based on satellite imagery. CNN's are starting to be used (although timidly) in this field. For instance, in [23], a CNN was used to process ground velocity records to detect earthquakes. Nonetheless, CNN's have been used to extract information and classify Google Street View images, including street number recognition [24], traffic sign recognition [25], number of building stories [26], and to estimate neighborhood demographic composition [27].

3.1 Network Architectures and Model Training

In this work, we consider five state-of-the-art CNN networks with different levels of architecture complexity (see Fig. 4). These networks have presented excellent performances on large scale classification challenges, such as Imagenet with over 1.5 Million images and 1000 classes [28]: VGG16, VGG19, InceptionV3, ResNet50 and Xception. In **Vgg16** and **Vgg19** images are passed through a stack of convolutional layers using filters with a very small receptive field, which reduces the number of parameters on the network, while increasing depth in the network. These networks have 16 and 19 weight layers respectively [29]. **InceptionV3** sacrifices the architecture simplicity from Vgg16/Vgg19 to obtain significant reductions in computational cost while managing state-of-the-art performance in image classification tasks. To achieve this goal the Inception module is introduced, which is a block of very small filters at different sizes to capture meaningful features at different scales. A sequence of convolutional layers followed by several Inception modules using different filter sizes comprise the 48 layers in the network [30]. As the number of layers in the network increases, they tend to present a vanishing gradient problem where earlier layers in the network learn in a slow-paced manner compared to later layers in the network. This problem increases the training time and diminishes the prediction accuracy of the model, **ResNet50** manages this by using shortcut connections between blocks of convolutional layers to allow gradient information being transmitted further down the network. This network has 50 weight layers, as the name suggests [31]. Finally, **Xception** builds on the previous architectures presented and



introduces a two-step convolution where a spatial convolution is performed independently for channel in the image, followed by a 1x1 convolution across all the channels. This network is 71 layers deep [32].

Since these architectures have been pretrained on Imagenet, we are able to use transfer-learning to improve generalization performance on our classification task [33]. This technique consists of reusing generic features learnt on earlier layers in the network, while retraining later layers to learn more specific filters related to our dataset.

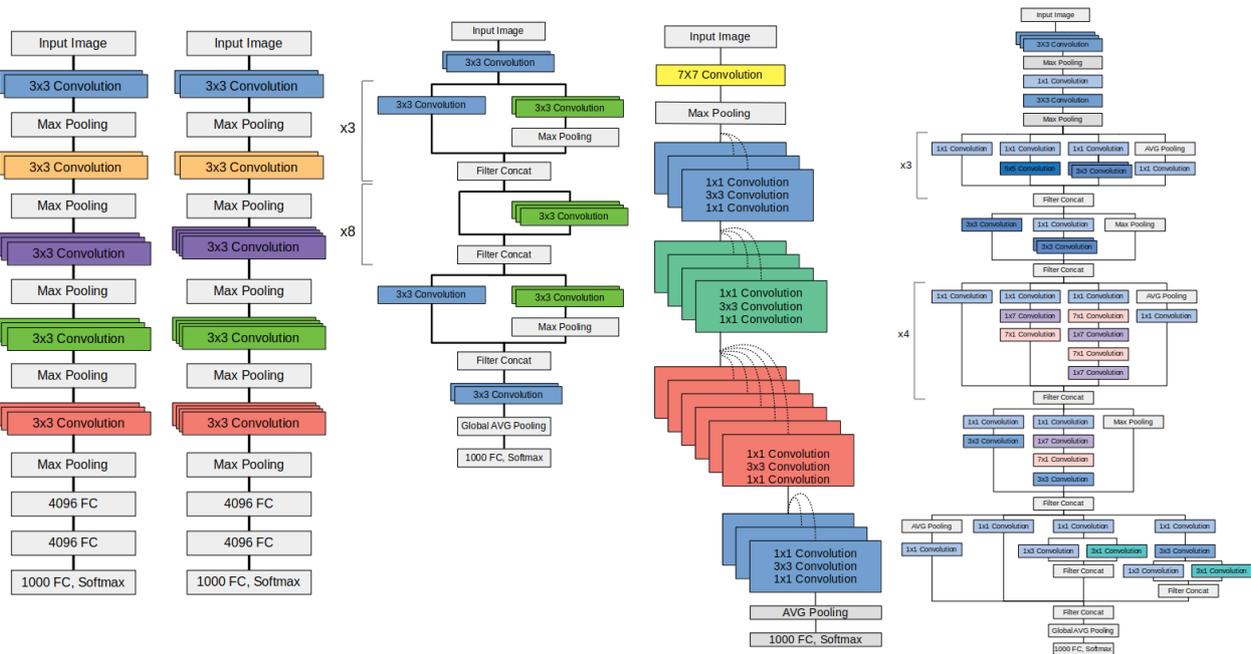


Fig. 4 – From left to right, Architectures of Vgg16, Vgg19, Xception, Resnet50 and InceptionV3. Adapted from [34]

3.2 Experimental methodology

Our aim is to develop a method capable of classifying a street level image of a building according to the presence of a rigid or flexible diaphragm structure. We believe that a good classifier, deployed under a suitable platform (i.e., a mobile phone) would reduce the cost and democratize risk assessment across wider population areas and segments. As in any predictive modelling task, care must be taken to devise the experimental methodology so that it is statistically sound and we do not incur in selection biases and reduce the risk of obtaining unrealistic performance estimates (too optimistic or too pessimistic).

Overall our process consists of five steps: (1) split the dataset; (2) define what performance metrics are most suitable to assess the models; (3) use specific dataset splits to train a variety of CNN architectures and select the best performing CNN; (4) use the remaining splits to obtain an unbiased estimate of the performance; and (5) inspect CNN internals as an attempt to understand the visual structures that most influence the CNNs in their decision.

We split our dataset in three parts: train (60%), validation (20%) and test (20%). This split is made in a stratified manner, such that the class proportions are maintained in all splits. We then train with fine tuning (see above) each one of the selected architectures using the train split and measure its performance using the train and validation splits. This performance is measured in terms of accuracy (percentage of correct model predictions), precision and recall (see next section), which overall allow us to assess how well models are



distinguishing both image classes (flexible and rigid). The validation split has not been seen by the training procedure and thus measures the generalization capabilities of each CNN model with unseen data. The difference in train and validation performances allows us to assess other aspects of the procedure itself, such as if the CNN models overfit (i.e., “memorized” part of the data instead of finding distinctive image features that can be used on new images). Finally, we select the best performing CNN model according to the validation split, but we use the test split to report performance. This avoids selection bias [35], this is, the data used to assess the overall process performance should not have been used in any processing stage, neither training nor model selection.

3.3 Performance metrics

In this paper we use three performance metrics: precision, recall and accuracy. For a better explanation of our performance measures, we will present the confusion matrix in Table 1. Precision is a measure of exactness calculated as the percentage of buildings classified by a model as having flexible diaphragm that were correct (i.e., $TP/(TP+FP)$). Recall is a measure of completeness calculated as the percentage of building with flexible diaphragm that were correctly identified as such (i.e., $TP/(TP+FN)$). Finally, accuracy is calculated as the percentage of buildings that were correctly classified (i.e., $(TP+TN)/(TP+TN+FP+FN)$). It is a widespread overall performance metric but must be used with care especially with imbalanced data. Observe that, in our case, we have 66% flexible diaphragms and 34% of rigid diaphragms, and that a dumb classifier that labels every image as flexible will attain 66% accuracy (100% in flexible and 0% in rigid). This would be very different from a 66% accuracy that shows evenly in both classes. This is the reason by which it must be interpreted with complementary metrics, such as precision and recall in this case.

Table 1 – Elements of a confusion matrix

Result	Actual Class: flexible	Actual Class: rigid
Predicted class: flexible	True Positive (TP)	False Positive (FP)
Predicted class: rigid	False Negative (FN)	True Negative (TN)

4. Analysis of results

Table 2 presents the performance metrics for the validation and test datasets. Recall that to avoid biases we use performance on the validation set to select which model to use, and performance on the test set to report the performance on the selected model and decision making. According to this we select Vgg19 as the best model, since it in validation has outstanding recall performance (0.83) and accuracy (0.84) with respect to the rest, and still precision is second best (0.92). For decision making (for instance, when assessing whether it is worthwhile or cost-effective to introduce the model in a production environment or application), we use the performance of the selected architecture on the test set (0.88 precision, 0.84 performance, 0.80 accuracy).

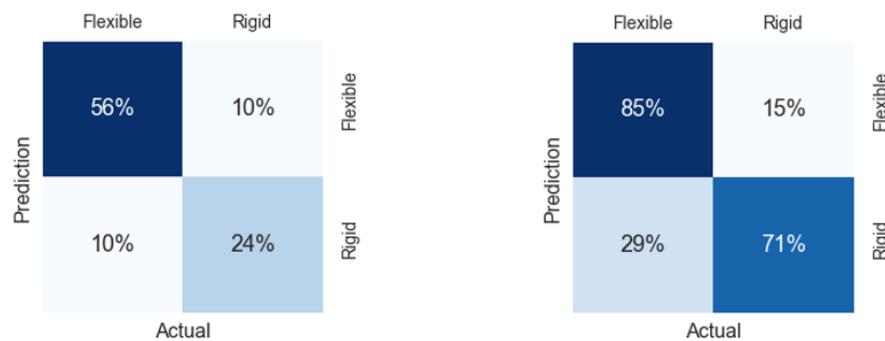
Figures 5a and 5b present two complementary versions of confusion matrices. Fig. 5a presents the overall magnitude of correct and incorrect predictions with respect to the total population (i.e., the entire matrix adds up to 100%). The results indicate that 80% of the images are correctly classified and the remaining 20% mistakenly classified images are evenly distributed between false positives and false negatives. Fig. 5b focuses on accuracy per class by presenting a version of the confusion matrix in which each row adds up to 100%. The results show that the accuracy level is considerably higher for flexible diaphragms than for rigid diaphragms, 85% vs. 71%, this difference may be a consequence of having a considerably higher number of buildings with flexible diaphragms than those with rigid diaphragm.



Table 2 – Performance metrics

Network	Validation Dataset			Test Dataset		
	Precision	Recall	Accuracy	Precision	Recall	Accuracy
Vgg16	0.97	0.68	0.77	0.92	0.73	0.77
Vgg19	0.92	0.83	0.84	0.88	0.84	0.80
InceptionV3	0.83	0.72	0.71	0.88	0.79	0.78
Xception	0.90	0.76	0.77	0.88	0.83	0.80
Resnet50	0.92	0.56	0.66	0.85	0.52	0.60

The similarity between the performance metrics obtained in the validation and test databases is also significant, as it signals the fact that both datasets are equally representative of our underlying buildings distribution and reduces the risk of selection bias.



a) Percentage based on total number of buildings

b) Percentage based on buildings of each class

Fig. 5 – Confusion matrix for selected architecture (Vgg19). a) Distribution based on the total number of buildings (all the values add to 100%), b) Distribution based on the building class (each row adds to 100%)

As we pointed out in Section 2, each expert assigned a certainty level of 100%, 75%, or 50% when classifying each image. On the other hand, in each classification the CNN assigns a probability of belonging to each of the possible categories (i.e., flexible, rigid), thus, an image is classified as flexible if the probability assigned to that category is the highest. Taking the category of flexible diaphragm as a reference, Fig. 6 presents a comparison between the level of certainty of the human vs. the probability assigned by CNN. Note that, although the level of certainty assigned by the expert is not taken into account in any part of the classification process, there is a direct relationship between the two measures when the probability of certainty of the expert is greater than or equal to 25%.

However, when the expert is confident about a rigid classification (left side of Fig. 6, the probability of flexible is close to 0%, which means a probability of rigid assigned by the expert close to 100%) our model is



most confused (leftmost blue bar in the figure around 40%). We believe this signals the fact that we have an unbalanced dataset with twice fewer rigid examples than flexible ones. Therefore, the model has not been given the chance to properly generalize what visual features are mostly associated to rigid structures and, thus, being more unsure when classifying an image as such.

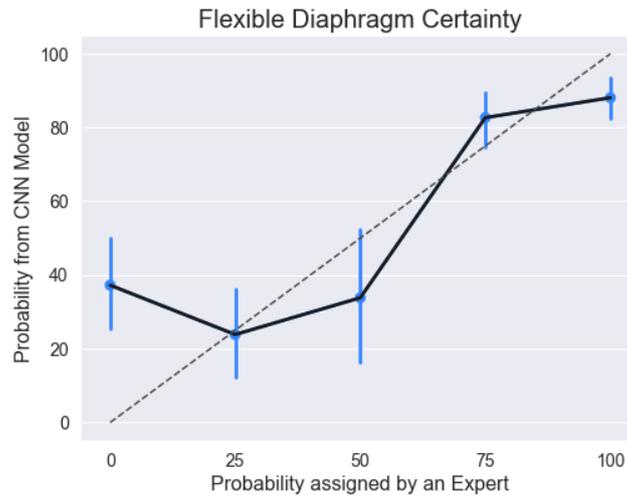


Fig. 6 – Comparison of the building probability of being flexible assigned by expert and the network

5. Conclusions

In this paper we explore the potential of CNN for classifying street-level imagery of one-story unreinforced masonry buildings according to the flexibility (rigid or flexible) of the roof diaphragm. For this exercise we used 1,122 images from the Metropolitan Area of Medellín. Each image was labeled by experts as flexible or rigid roof diaphragm.

The results show that Vgg19 is the best architecture for the problem of diaphragm classification, with an accuracy of 0.80, a precision of 0.88, and a recall of 0.84. According to the confusion matrices, 80% of the images are correctly classified and the accuracy level is considerably higher for flexible diaphragms than for rigid diaphragms, 85% vs. 71%. We also found that there is a direct relationship between the level of certainty assigned by the experts when classifying an image and the probability assigned by the CNN. This highly automated method stands out as a good candidate to offer a fast and inexpensive way to build an exposure model for seismic risk assessment of one-story unreinforced masonry buildings. Moreover, we believe that more accurate results could be obtained if additional information from satellite images (such as roof type) are considered into the building classification. We will explore this possibility in further analyses.

Results show that the identification of the buildings with rigid diaphragm is not as accurate as the one for buildings with flexible diaphragm. Upcoming work will concentrate on increasing the number of images with rigid diaphragm in the database in order to improve the machine learning process for such building category. In addition, we believe that the learning process can be improved if images with external objects such as vegetation, signs, vehicles, etc. are not included in the training process. For such purpose, a process could be implemented to remove objects that may generate a distraction to the machine. Another option worth exploring is the use of eye-tracking devices to identify those areas of the image considered by the expert when classifying an image and then fit this information into the CNN. This will allow minimizing the influence of external objects without having to remove them.

We are encouraged by the results to extend the methodology to buildings with two or more number of stories. For such buildings an interesting feature that could also be considered in addition to the identification of the flexibility of the floor(s) and roof diaphragm is the presence of different lateral load-resisting systems



in the building, a situation that takes place due to informal construction processes in which stories are added to existing buildings without following construction best practices.

6. Acknowledgements

We want to extend our gratitude to the Apolo scientific computing center at EAFIT University that allows us to significantly reduce the computational process involved in this research.

7. References

- [1] Silva V, Crowley H, Jaiswal K, Acevedo AB, Pittore M, Journey M (2018): Developing a global earthquake risk model. *16th European Conference on Earthquake Engineering*, Thessaloniki, Greece.
- [2] FEMA (2003): Inventory data: Collection and classification, Chapter 3, *HAZUS-MH 2.1 Technical Manual*, Washington DC, USA.
- [3] Acevedo AB, Jaramillo JD, Yepes C, Silva V, Osorio FA, Villar M (2017): Evaluation of the seismic risk of the unreinforced masonry building stock in Antioquia, Colombia. *Natural Hazards*, **86**, 31-54.
- [4] Yepes-Estrada C, Silva V, Valcárcel J, Acevedo AB, Tarque N, Hube MA, Coronel G, Santa María H (2017): Modeling the residential building inventory in South America for seismic risk assessment. *Earthquake Spectra*, **33** (1), 299-322.
- [5] Salgado-Gálvez MA, Zuloaga-Romero D, Bernal GA, Mora MG, Cardona OD (2014): Fully probabilistic seismic risk assessment considering local site effects for the portfolio of buildings in Medellín, Colombia. *Bulletin of Earthquake Engineering*, **12**, 671-695.
- [6] Herrera DC, Moreno JD, Garzón D, Daza JM, Yamín LE (2019): Evaluación del riesgo sísmico de edificaciones en el Área Metropolitana del Valle de Aburrá. *IX Congreso Nacional de Ingeniería Sísmica*, Cali, Colombia.
- [7] Fernández RI, González R, García A, Yamín LE (2019): Funciones de vulnerabilidad sísmica para mampostería no reforzada. *IX Congreso Nacional de Ingeniería Sísmica*, Cali, Colombia.
- [8] Santa María H, Hube MA, Rivera F, Yepes-Estrada C, Valcárcel J (2017): Development of national and local exposure models of residential structures in Chile. *Natural Hazards*, **86**, 55-79.
- [9] Bengio Y, Courville A, Vincent P (2013): Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, **35**, 1798–1828.
- [10] Khan A, Sohail A, Zahoor U, Qureshi AS (2019): A survey of the recent architectures of deep convolutional neural networks. *arXiv preprint arXiv:1901.06032*.
- [11] Caputo AC, Vigna A (2017): Numerical simulation of seismic risk and loss propagation effects in process plants: An oil refinery case study. *ASME 2017 Pressure Vessels and Piping Conference*. American Society of Mechanical Engineers Digital Collection.
- [12] Esposito S, Iervolino I, d'Onofrio A, Santo A, Cavalieri F, Franchin P (2015): Simulation-based seismic risk assessment of gas distribution networks. *Computer-Aided Civil and Infrastructure Engineering*, **30** (7), 508-523.
- [13] Cook D, Fitzgerald K, Chrupalo T, Haselton CB (2017): Comparison of FEMA P-58 with other building seismic risk assessment methods. Retrieved from http://static-assets.hbrisk.com/reports/RiskAssessmentComparisonReport_FEMA-P-58-to-OtherMethods_2017_03_08.pdf.
- [14] Frolova NI, Larionov VI, Bonnin J, Sushchev S P, Ugarov AN, Kozlov, MA (2017): Seismic risk assessment and mapping at different levels. *Natural Hazards*, **88** (1), 43-62.
- [15] Mangalathu S, Jeon JS, Padgett JE, DesRoches R (2017): Performance-based grouping methods of bridge classes for regional seismic risk assessment: Application of ANOVA, ANCOVA, and non-parametric approaches. *Earthquake Engineering & Structural Dynamics*, **46** (14), 2587-2602.
- [16] Riedel I, Guéguen P, Dalla Mura M, Pathier E, Leduc T, Chanussot J (2015): Seismic vulnerability assessment of urban environments in moderate-to-low seismic hazard regions using association rule learning and support vector machine methods. *Natural hazards*, **76** (2), 1111-1141.



- [17] Zhang Y, Burton H., Sun H, Shokrabadi M (2018): A machine learning framework for assessing post-earthquake structural safety. *Structural Safety*, **72**, 1-16.
- [18] Fisher WD, Camp TK, Krzhizhanovskaya VV (2017): Anomaly detection in earth dam and levee passive seismic data using support vector machines and automatic feature selection. *Journal of Computational Science*, **20**, 143-153.
- [19] Bai Y, Adriano B, Mas E, Gokon H, Koshimura S (2017): Object-based building damage assessment methodology using only post event ALOS-2/PALSAR-2 dual polarimetric SAR intensity images. *Journal of Disaster Research*, **12** (2), 259-271.
- [20] Wieland M, Pittore M, Parolai S, Zschau J (2012): Exposure estimation from multi-resolution optical satellite imagery for seismic risk assessment. *ISPRS International Journal of Geo-Information*, **1** (1), 69-88.
- [21] Geiß C, Schauß A, Riedlinger T, Dech S, Zelaya C, Guzmán N, Taubenböck H (2017): Joint use of remote sensing data and volunteered geographic information for exposure estimation: evidence from Valparaíso, Chile. *Natural Hazards*, **86** (1), 81-105.
- [22] Raouf A, Peng Y, Shah T (2017): Integrated use of aerial photographs and LiDAR images for landslide and soil erosion analysis: a case study of Wakamow Valley, Moose Jaw, Canada. *Urban Science*, **1** (2), 20.
- [23] Perol T, Gharbi M, Denolle M (2018): Convolutional neural network for earthquake detection and location. *Science Advances*, **4** (2), e1700578.
- [24] Goodfellow IJ, Bulatov Y, Ibarz J, Arnoud S, Shet V (2013): Multi-digit number recognition from street view imagery using deep convolutional neural networks. *arXiv preprint*, arXiv:1312.6082.
- [25] Jin J, Fu K, Zhang C (2014): Traffic sign recognition with hinge loss trained convolutional neural networks. *IEEE Transactions on Intelligent Transportation Systems*, **15** (5), 1991-2000.
- [26] Iannelli G, Dell'Acqua F (2017): Extensive exposure mapping in urban areas through deep analysis of street-level pictures for floor count determination. *Urban Science*, **1**(2), 16.
- [27] Gebru T, Krause J, Wang Y, Chen D, Deng J, Aiden EL, Fei-Fei L (2017): Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States. *Proceedings of the National Academy of Sciences*, **114**(50), 13108-13113.
- [28] Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, Berg AC, Fei-Fei L (2015): Imagenet large scale visual recognition challenge, *International Journal of Computer Vision*, **115**, 211–252.
- [29] Simonyan K, Zisserman A (2014): Very deep convolutional networks for large-scale image recognition, *arXiv preprint*, arXiv:1409-1556.
- [30] Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z (2016): Rethinking the inception architecture for computer vision. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2818–2826.
- [31] He K, Zhang X, Ren S, Sun J (2016): Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- [32] Chollet F, (2017): Xception: Deep learning with depthwise separable convolutions. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1251–1258.
- [33] Yosinski J, Clune J, Bengio Y, Lipson H (2014): How transferable are features in deep neural networks?, *Advances in neural information processing systems*, 3320–3328.
- [34] Leonardo MM, Carvalho TJ, Rezende E, Zucchi R, Faria FA (2018): Deep feature-based classifiers for fruit fly identification (diptera: Tephritidae). In *2018 31st SIBGRAPI Conference on Graphics, Patterns and Images*, 41-47.
- [35] Kuhn M, Johnson K (2013): *Applied predictive modeling*. Springer-Verlag New York, USA.