Predicting future myocardial infarction from angiographies with deep learning

Dorina Thanou[1] **Ortal Senouf**[1] **Omar Raita**[1] **Emmanuel Abbé**[1] **Pascal Frossard**[1]
Farhang Aminfar[2] **Denise Auberson**[2] **Nicolas Dayer**[2] **David Meier**[2] **Mattia Pagnoni**[2]
Olivier Muller[2] **Stéphane Fournier**[2] **Thabo Mahendiran**[2]
[1]EPFL, Switzerland, [2]CHUV, Switzerland
{dorina.thanou, ortal.senouf, omar.raita, emmanuel.abbe, pascal.frossard}@epfl.ch
{farhang.aminfar, denise.auberson, david.meier, olivier.muller}@chuv.ch
mattiapagnoni91@gmail.com, {stephane.fournier, thabodhan.mahendiran}@chuv.ch

Abstract

In patients with stable Coronary Artery Disease (CAD), the identification of lesions which will be responsible of a myocardial infarction (MI) during follow-up remains a daily challenge. In this work, we propose to predict culprit stenosis by applying a deep learning framework on image stenosis obtained from patient data. Preliminary results on a data set of 746 lesions obtained from angiographies confirm that deep learning can indeed achieve significant predictive performance, and even outperforms the one achieved by a group of interventional cardiologists. To the best of our knowledge, this is the first work that leverages the power of deep learning to predict MI from coronary angiograms, and it opens new doors towards predicting MI using data-driven algorithms.

1 Introduction

Coronary Artery Disease (CAD), a leading cause of death worldwide, involves narrowing (i.e., stenosis) of at least one of the large arteries supplying the heart. In patients with stable CAD, the identification of lesions which will be responsible of a myocardial infarction in the future, remains a daily challenge. Invasive coronary angiography (CA) is the gold standard for determining the presence and severity of coronary stenosis, which involves continuous X-ray (i.e., fluoroscopy) with simultaneous injection of radiopaque contrast into the coronary arteries. The identification of a future culprit lesion, i.e., a lesion or stenosis that will lead to an MI is particularly challenging as these lesions are generally not fully characterized. In current clinical practice, stenosis severity is still often determined by physician visual assessment only. With this approach, a 70% diameter stenosis is generally considered an indicator of clinically significant lesions. This approach, while considered as clinical standard, has known limitations, such as significant intra- and inter-rater variability, as well as high positive prediction bias.

At the same time, deep learning has recently been successful in medical applications where the data is abundant and the labels are highly reliable such as melanoma (Esteva et al., 2017), and glaucoma (Chen et al., 2015) detection. However, the heart attack prediction problem has several new layers of complexities: (i) The heart and coronaries are internal and mobile structures, requiring imaging techniques that are more complex and less scalable than skin or retina pictures, and resulting in smaller and noisier datasets. (ii) Angiographic images consists of 2D images representing the 3D anatomy, which leads to some loss of information. (iii) CAD is a diffuse disease process, which results from a complex interplay of local and systemic factors ranging from vascular wall histology to fluid hemodynamics, and which in turn lead to complex target functions. We note that deep learning has already been applied to coronary angiograms, with however main focus on stenosis detection

Medical Imaging Meets NeurIPS workshop (NeurIPS 2021), Sydney, Australia.

(Ovalle-Magallanes et al., 2020; Au et al., 2018; Zhou et al., 2021; Du et al., 2020; Moon et al., 2021; Danilov et al., 2021) or vessel segmentation (Yang et al., 2019; Jun et al., 2020; Wang et al., 2020), but not MI prediction.

The aim of this preliminary work is to prove that machine learning can contribute towards predicting MI, even with scarce and imperfect data. In particular, we apply state-of-the-art deep learning techniques on coronary angiograms obtained from patients who had an MI in a time frame of 5 years after the angiogram had been taken. This way, it was possible to annotate which lesions, seen in the angiograms, lead to a future MI. Each patient presented with an MI, contains stenosis across different locations of the coronary tree. We study the problem at a stenosis level, i.e., we treat each stenosis independently as a patch that is given as an input to a deep learning architecture (see Fig. 1). The network, which is built on a ResNet18 (He et al., 2016), is trained on patches that are labeled as culprit and non-culprit. To compensate for the small size of the dataset, we perform data augmentation in the training phase. The obtained results are quite encouraging as they achieve a non-trivial prediction performance of approximately 80%. This level of performance is higher with respect to the one achieved by interventional cardiologists. We emphasize that the limited amount of images used for training as well as the unbalanced nature of the dataset makes the specific setup very challenging for deep learning. We thus believe that, with a larger dataset, the current work can provide a different perspective in treating coronary artery diseases. It will hopefully generate more research towards leveraging the power of machine learning, and in particular extending existing algorithms, to address the challenges posed by the heart attack prediction problem.

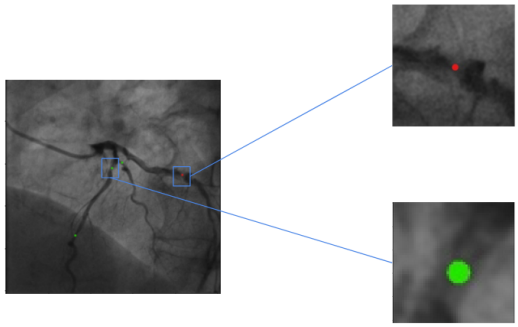

Figure 1: Examples of patches (on the right) extracted from an annotated CA (on the left). Red and green color represents culprit and non-culprit lesions respectively.

2 Methodology and Results

2.1 Clinical study

Our analysis is based on a retrospective study including 83 patients, from the Lausanne University Hospital (CHUV). For each patient, one to four image frames were manually selected from the clinicians out of the entire DICOM video sequence. The raw data is composed of 374 anonymized and labelled images of 1014x1014 pixels each, that have been manually annotated at a stenosis level by interventional cardiologists. The final dataset is constructed by extracting the corresponding patches of 224x224 pixels centered around the stenotic region from the raw images, in such a way that only one stenosis is included per patch. A label is assigned to each patch i.e., 1 for culprit and 0 for non-culprit (see Fig. 1). The total number of patches is 746, out of which only 203 are culprit ($\approx 27\%$).

2.2 Deep learning framework and results

With these patches, we train a state-of-the-art ResNet18 (He et al., 2016). A block diagram containing the building blocks of such an architecture is shown in Fig. 2. In particular, out of all patches, 80% are used for training, 10% for validation, and 10% for testing the deep learning algorithm. Due to the very unbalanced nature of the dataset, we augment the training dataset by randomly replicating some of the culprit patches, in such a way that we obtain a balanced dataset i.e., 439 culprit and 439

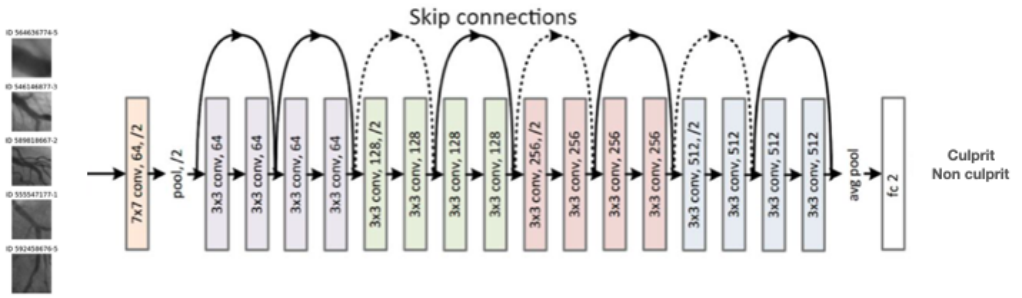

Figure 2: Building blocks of the ResNet architecture. The input is processed through 18 stages; many of these stages include connections that transmit the input to a later stage through a skip connection. The final stage is a fully connected layer that provides a classification decision: culprit or non-culprit. The text in each of the block describes its key properties: N x N conv is the kernel size; the next integer is the number of kernels; if present /N describes the spatial sub-sampling (stride).

non-culprit lesions. Then, we apply a sequence of transformations from the Albumentations library[1], that consists of i) Median Blur, ii) Rotation, iii) ShiftScaleRotate, iv) Resize to 224x224. Note that in order to highlight some important structures in the images, blurring is applied to validation and testing set as well.

We initialize the network using a pretrained model from ImageNet. The performance of each network is measured using four evaluation metrics: (i) classification accuracy, which measures the performance of correctly predicted classes; (ii) sensitivity which measures the proportion of future culprit lesions that are correctly identified; (iii) specificity which measures the proportion of non-culprit lesions that are correctly identified; and (iv) F1 score, which is defined as the harmonic mean of precision and recall (Manning et al., 2008). The hyper-parameters of the network, i.e., the learning rate, and the weight decay, are determined using 10-fold cross validation on the training and validation set, and they are set to 0.0015 and 0.1113 respectively. These hyperparameters are chosen based on the maximum F1 score. With these hyperparameters, we train a network, by minimizing the binary cross entropy loss with a stochastic gradient descent optimizer, using both training and validation data. The batch size is fixed to 20. The test set is used only for the final evaluation of the model. Due to the very limited number of patients, training is performed patch-wise, i.e., we do not take into account the correspondence between different views of the same stenosis. However, since the goal is the classification of a stenosis, in the test phase, we predict a stenosis as culprit if and only if at least a patch coming from one of the views is classified as culprit. This approach is closer to how cardiologists evaluate a stenosis in clinical settings. For comparison, the exact set of testing lesions is evaluated by two blinded interventional cardiologists.

A comparison between the prediction of the deep learning algorithm and the human prediction is shown in Table 1. We observe that our learning scheme outperforms the prediction of both clinical experts with respect to all metrics. Although the impact of this study is limited by the limited number of available patients, and the corresponding stenosis, the non-trivial performance obtained with the current data is promising, and it indicates that deep learning can have the potential to become a game changer in hard problems in medicine, such as the one studied in this paper. We thus believe that this line of research could open new perspectives in understanding the causes of MI, and eventually reach clinical translation. In that direction, we plan to extend the existing work by looking into aspects such as the interpretability of the results, as well as introducing prior knowledge in the learning algorithm by taking into account the hemodynamics. We reiterate that the current study is at a lesion level. Extension to a patient level study is definitely worth investigating, and would allow us to integrate other clinical aspects in the framework.

[1] `https://albumentations.ai/docs/`

Table 1: Performance evaluation

	Classification accuracy	Sensitivity	Specificity	F1 score
Deep learning	**0.775**	**0.667**	**0.806**	**0.571**
Cardiologist 1	0.625	0.444	0.677	0.348
Cardiologist 2	0.575	0.222	0.677	0.19

References

Au, B., Shaham, U., Dhruva, S., Bouras, G., Cristea, E., MD, A. L., Coppi, A., Warner, F., Li, S.-X., and Krumholz, H. (2018). Automated characterization of stenosis in invasive coronary angiography images with convolutional neural networks. *arXiv:1807.10597*.

Chen, X., Xu, Y., Kee Wong, D. W., Wong, T. Y., and Liu, J. (2015). Glaucoma detection based on deep convolutional neural network. In *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 715–718.

Danilov, V., Klyshnikov, K., OM, G., Kutikhin, A., Ganyukov, V., Frangi, A., and Ovcharenko, E. (2021). Real-time coronary artery stenosis detection based on modern neural networks. *Nature Scientific Report*, 11(1).

Du, T., Xie, L., Zhang, H., Liu, X., Wang, X., Chen, D., Xu, Y., Sun, Z., Zhou, W., Song, L., Guan, C., Lansky, A., and Xu, B. (2020). Automatic and multimodal analysis for coronary angiography: training and validation of a deep learning architecture. *EuroIntervention*.

Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., and Thrun, S. (2017). Dermatologist-level classification of skin cancer. *Nature*, 39:1545–1557.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.

Jun, T. J., Kweon, J., Kim, Y.-H., and Kim, D. (2020). T-net: Nested encoder–decoder architecture for the main vessel segmentation in coronary angiography. *Neural Networks*, 128:216–233.

Manning, C. D., Raghavan, P., and Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press.

Moon, J., Lee, D., Cha, W., Chung, M., Lee, K., Cho, B., and Choi, J. (2021). Automatic stenosis recognition from coronary angiography using convolutional neural networks. *Comput Methods Programs Biomed*.

Ovalle-Magallanes, E., Avina-Cervantes, J. G., Cruz-Aceves, I., and Ruiz-Pinales, J. (2020). Transfer learning for stenosis detection in x-ray coronary angiography. *Mathematics*, 8(9).

Wang, W., Wang, H., Chen, Q., Zhou, Z., Wang, R., Wang, H., Zhang, N., Chen, Y., Sun, Z., and Xu, L. (2020). Coronary artery calcium score quantification using a deep-learning algorithm. *Clinical Radiology*, 75(3).

Yang, S., Kweon, J., Roh, J. H., Lee, J. H., Kang, H., Park, L. J., Kim, D. J., Yang, H., Hur, J., Kang, D. Y., Lee, P. H., Ahn, J. M., Kang, S. J., Park, D. W., Lee, S. W., Kim, Y. H., Lee, C. W., Park, S. W., and Park, S. J. (2019). Deep learning segmentation of major vessels in x-ray coronary angiography. *Nature Scientific reports*, 9(1).

Zhou, C., Dinh, T. V., Kong, H., Yap, J., Yeo, K. K., Lee, H. K., and Liang, K. (2021). Automated deep learning analysis of angiography video sequences for coronary artery disease. *arXiv:2101.12505*.