

# AUTOMATED 5-YEAR MORTALITY PREDICTION USING DEEP LEARNING AND RADIOMICS FEATURES FROM CHEST COMPUTED TOMOGRAPHY

Gustavo Carneiro<sup>◇</sup>, Luke Oakden-Rayner<sup>†</sup>, Andrew P. Bradley<sup>•</sup>, Jacinto Nascimento<sup>\*</sup>, Lyle Palmer<sup>† \*</sup>

<sup>◇</sup> Australian Centre for Visual Technologies, The University of Adelaide, Australia

<sup>†</sup> School of Public Health, The University of Adelaide, Australia

<sup>•</sup> School of ITEE, The University of Queensland, Australia

<sup>\*</sup> Institute for Systems and Robotics, Instituto Superior Técnico, Portugal

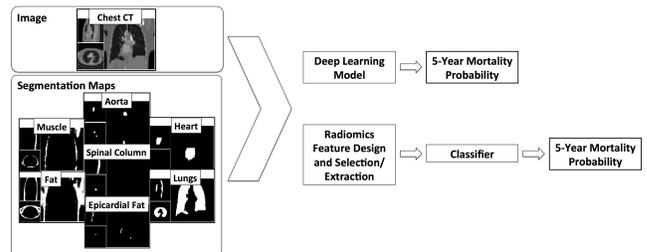
## ABSTRACT

In this paper, we propose new prognostic methods that predict 5-year mortality in elderly individuals using chest computed tomography (CT). The methods consist of a classifier that performs this prediction using a set of features extracted from the CT image and segmentation maps of multiple anatomic structures. We explore two approaches: 1) a unified framework based on two state-of-the-art deep learning models extended to 3-D inputs, where features and classifier are automatically learned in a single optimisation process; and 2) a multi-stage framework based on the design and selection and extraction of hand-crafted radiomics features, followed by the classifier learning process. Experimental results, based on a dataset of 48 annotated chest CTs, show that the deep learning models produces a mean 5-year mortality prediction AUC in [68.8%,69.8%] and accuracy in [64.5%,66.5%], while radiomics produces a mean AUC of 64.6% and accuracy of 64.6%. The successful development of the proposed models has the potential to make a profound impact in preventive and personalised healthcare.

**Index Terms**— deep learning, radiomics, feature learning, hand-designed features, computed tomography, five-year mortality

## 1. INTRODUCTION

Cause of death is a complex question because ill-health and comorbidities strongly influence mortality, but may not be listed as the primary diagnosis - in fact many significant diseases are never diagnosed at all. We believe that the medical imaging community should focus on tools to improve prognosis given that diagnosis is variable and does not capture the range of human health, whereas outcomes like death are unequivocal and better reflect the underlying status of the body. In addition, the prediction of reduced life expectancy in individuals is a public health priority and central to personalised medical decision making [1]. Previous attempts to predict reduced life expectancy in the elderly have been studied using invasive (e.g., blood samples) and non-invasive (e.g., self-reported survey results, clinical examination) tests [1]. These



**Fig. 1.** The proposed deep learning and radiomics models use image and segmentation maps to estimate the patient’s 5-year mortality probability.

approaches resulted in a classification accuracy between 60% and 80% [1, 2], although patient age alone has shown a predictive accuracy of above 65% [1]. Compared to these previous attempts, the use of chest CT for the prediction of reduced life expectancy is advantageous because these scans potentially offer information on multiple organs and tissues from a single non-invasive test. Hence, it is the aim of this paper to show that the use of chest CT alone (i.e., excluding previously used invasive and demographic markers such as age and gender) can produce accurate prediction of reduced life expectancy.

Typically, prognostic models in medical image analysis have been designed for the prediction of disease specific outcomes [3, 4, 5, 6, 7], where the methodology requires hand-crafted features. These features are selected/extracted based on their correlation with the prognosis, followed by modelling of the desired outcome using survival models or predictive classifiers. This multi-stage process of feature design and selection/extraction, followed by modelling has many disadvantages, such as the hand-crafting of the image features requiring medical expertise and being useful only for the particular prognosis being addressed, and the independence between feature selection/extraction and modelling potentially introducing redundant features and removing complementary features for the classification process. Furthermore, recent studies in this field do not address the problem of predicting more general life expectancy in individuals from chest CTs.

In this paper, we propose new approaches for the prediction of 5-year all-cause mortality in elderly individuals

\*This work was partially supported by the Australian Research Council’s Discovery Projects funding scheme (project DP140102794). Prof. Bradley is the recipient of an Australian Research Council Future Fellowship(FT110100623).

using chest CT and the segmentation maps of the following anatomies: aorta, spinal column, epicardial fat, body fat, heart, lungs and muscle. We have chosen chest CTs because they are commonly performed and widely available from hospitals, which facilitates dataset acquisition, and the segmentation maps are informed by previous biomarker research, which has demonstrated predictive and detectable changes in these tissues [5, 6, 7]. The approaches developed in the paper are the following (Fig. 1): 1) a unified framework based on two state-of-the deep learning models extended to 3-D inputs, where features and classifiers are automatically learned in a single optimisation; and 2) a multi-stage framework based on the hand-crafting and selection/extraction of radiomics features, followed by a classifier learning process. Cross-validated experiments based on 48 annotated chest CT volumes show that the deep learning model produces mean classification AUC in [68.8%,69.8%] and accuracy in [64.5%,66.5%], while radiomics produces a mean AUC of 64.6% and accuracy of 64.6%. Note that this is currently the largest such dataset in the field (e.g., the most similar dataset in the field is Visceral Anatomy 3, containing only 20 CT volumes<sup>1</sup>). Even though these results show comparable classification accuracy, deep learning models have an important advantage compared to radiomics: the fully automated approach to designing features, without requiring the assistance of a medical expert.

## 2. RELATED WORK

This paper is related to radiomics and deep learning for medical image analysis. Radiomics methods are concerned with the design of hand-crafted features and their association with subtle variations in disease processes [4]. Usually, radiomics methods are applied to imaging studies of patients with active tumours [3], but the application of these techniques to a general population of radiology patients for the prediction of important medical outcomes (e.g., mortality) is novel. In this application such hand-crafting of features is inefficient because it requires medical expertise, or alternatively if the features are task-agnostic (i.e. not informed by domain knowledge) it is not possible to know in advance which features will be effective, and it is therefore necessary to generate many possible features. This often requires a separate feature selection/extraction step to reduce the training complexity of the final classifier, and often is based on a search heuristic that is not necessarily linked to the classification target. For every new problem being addressed by radiomics, these two sub-optimal steps must be repeated, representing the major disadvantage of these methods.

Deep learning models are defined by a network composed of several layers of non-linear transformations that represent features of different levels of abstraction extracted directly from the input data [8, 9, 10]. In medical image analysis, deep learning can significantly improve segmentation and classification results [11, 12, 13], but its application to routinely collected medical images to predict important medical outcomes (e.g., mortality) has yet to be demonstrated. Our main references are the multi-view classification of mammograms [14] and the chest pathology classification using X-Rays [11] be-

cause these works use deep learning methods for the high-level classification of medical images, but both classify a diagnostic outcome, which is conceptually different compared to our prognostic output.

## 3. METHODOLOGY

The **dataset** is represented by  $\mathcal{D} = \{(\mathbf{v}, \{\mathbf{s}^{(j)}\}_{j \in \mathcal{A}}, y)\}_{i=1}^{|\mathcal{D}|}$ , where  $\mathbf{V} : \Omega \rightarrow \mathbb{R}$  denotes the chest CT with  $\Omega \in \mathbb{R}^3$  representing the volume lattice of size  $w \times h \times d$ ,  $\mathbf{s}^{(j)} : \Omega \rightarrow \{0, +1\}$  represents the segmentation map for the anatomies in  $\mathcal{A} = \{muscle, body\ fat, aorta, spinal\ column, epicardial\ fat, heart, and\ lungs\}$ , and  $y \in \{0, 1\}$  denotes whether the patient is dead ( $y = 1$ ) or alive ( $y = 0$ ) on the time to censoring (time to death or time of last follow-up).

**Radiomics** approaches comprise the following stages [5]: 1) hand-crafting a large pool of features, 2) feature selection/extraction, and 3) classifier training. The hand-crafting process involves medical expertise to extract intensity, texture and shape information from particular image regions that are relevant for the final prognosis/diagnosis task. The feature extraction is denoted by

$$\mathbf{r} = r(\mathbf{v}, \{\mathbf{s}^{(j)}\}_{j \in \mathcal{A}}), \quad (1)$$

where  $r(\cdot)$  represents a function that extracts the features  $\mathbf{r} \in \mathbb{R}^R$ . Intensity features are based on the histogram of grey values  $\mathbf{h}^{(j)} \in \mathbb{R}^H$  per anatomy  $j \in \mathcal{A}$ . The feature is defined by statistics from  $\mathbf{h}^{(j)}$ , such as mean, median, range, skewness, kurtosis, and etc. In addition to these task-agnostic intensity-based features, we also include task-specific features that are related to the problem of estimating chronic disease burden, such as approximations of bone mineral density scoring (BMD) [6], emphysema scoring [7], and coronary (and aortic) artery calcification score [15].

The texture-based features use first and second-order matrix statistics, like the grey level co-occurrence matrix (GLCM) for anatomy ( $j$ ), denoted by  $\mathbf{M}_{GLCM}^{(j),d,a}$ , where the  $r^{th}$  row and  $c^{th}$  column of represent the number of times that grey levels  $r$  and  $c$  co-occur in two voxels separated by the distance  $d \in \mathbb{R}$  in the direction  $a \in \mathbb{R}$  within the segmentation map provided by  $\mathbf{s}^{(j)}$ . The grey level run-length matrix (GLRLM) for anatomy ( $j$ ) is defined by  $\mathbf{M}_{GLRLM}^{(j),a}$ , where the  $r^{th}$  row and  $c^{th}$  column denote the number of times a run of length  $c$  happens with grey level  $r$  in direction  $a$  within the segmentation  $\mathbf{s}^{(j)}$ . The grey level size-zone matrix (GLSZM) for anatomy ( $j$ ) is represented by  $\mathbf{M}_{GLSZM}^{(j)}$ , where the  $r^{th}$  row and  $c^{th}$  column denote the number of times  $c$  grey levels  $r$  are contiguous in 8-connected pixels within the segmentation  $\mathbf{s}^{(j)}$ . Finally, the multiple gray level size-zone matrix (MGLSZM) for anatomy ( $j$ ) is defined by  $\mathbf{M}_{MGLSZM}^{(j)}$ , computed by a weighted average of several  $\mathbf{M}_{GLSZM}^{(j)}$ , each estimated with a different number of possible grey levels. The features computed from these matrices are based on several statistics, such as energy, mean, entropy, variance, kurtosis, skewness, correlation, etc. Each of the intensity and texture features are defined in a spatial context, by the use of weighted mean positions and spatial quartile means

<sup>1</sup><http://www.visceral.eu/benchmarks/anatomy3-open/>

in all three dimensions, to identify any local variations across the tissues and organs. Finally, the shape-based features are based on the volume of each anatomy  $j \in \mathcal{A}$ , computed from the segmentation map  $\mathbf{s}^{(j)}$  [5]. The vector  $\mathbf{r}$  formed by such features (note that there may be a feature selection step to reduce the dimensionality of this vector) is used for training the classifier, as in:

$$\gamma^* = \arg \min_{\gamma} \sum_{i \in \mathcal{T}} \Delta_{radiomics}(y_i, g(\mathbf{r}_i; \gamma)), \quad (2)$$

where  $\mathcal{T} \in \mathcal{D}$  represents the training set,  $g(\mathbf{r}_i; \gamma)$  denotes a classifier that returns a value in  $[0, 1]$  indicating the confidence in the 5-year mortality prediction,  $\gamma$  represents the classifier parameters, and  $\Delta_{radiomics}(\cdot)$  denotes the loss function that penalises classification errors.

The **deep learning model** used in this work is the Convolutional Neural Network (ConvNet) [16, 9, 10], defined as follows:

$$f([\mathbf{v}, \{\mathbf{s}^{(j)}\}_{j \in \mathcal{A}}]; \theta) = f_{out} \circ f_L \circ \dots \circ f_2 \circ f_1([\mathbf{v}, \{\mathbf{s}^{(j)}\}_{j \in \mathcal{A}}]; \theta_1), \quad (3)$$

where  $\circ$  denotes the composition operator,  $\theta$  represents the ConvNet parameters (i.e., weights and biases), and the output is a value in  $[0, 1]$  indicating the confidence in the 5-year mortality prediction. Each network layer in (3) contains a set of filters, with each filter being defined by

$$\mathbf{x}(l+1) = f_l(\mathbf{x}(l); \theta_l) = \sigma(\mathbf{W}_l^\top \mathbf{x}(l) + \beta_l), \quad (4)$$

where  $\sigma(\cdot)$  represents a non-linearity [16],  $\mathbf{W}_l$  and  $\beta_l$  denote the weight and bias parameters, and  $\mathbf{x}(1) = [\mathbf{v}, \{\mathbf{s}^{(j)}\}_{j \in \mathcal{A}}]$ . The last layer  $L$  of the model in (3) produces a response  $\mathbf{x}(L+1)$ , which is the input for  $f_{out}(\cdot)$  that contains two output nodes (denoting the probability of 5-year mortality or survival), where layers  $L$  and  $out$  are fully-connected. The training of the model in (3) minimises the binary cross entropy loss on the training set  $\mathcal{T}$ , as follows:

$$\theta^* = \arg \min_{\theta} \sum_{i \in \mathcal{T}} \Delta_{conv}(y_i, f(\mathbf{x}_i(1); \theta)), \quad (5)$$

where  $\Delta_{conv}(y_i, f(\mathbf{x}_i(1); \theta)) = -y_i \times \log(f(\mathbf{x}_i(1); \theta)) - (1 - y_i) \times \log(1 - f(\mathbf{x}_i(1); \theta))$ .

## 4. EXPERIMENTS

**Materials and Methods:** The dataset has 24 cases (mortality) and 24 matched controls (survival), forming 48 annotated chest CTs of size  $512 \times 512 \times 45$ . Inclusion criteria for the mortality cases are: age  $> 60$ , mortality in 2014, and underwent CT chest imaging in the 3 to 5 years preceding death. Exclusion criteria are: acute disease identified on CT chest, mortality unrelated to chronic disease (e.g., trauma), and active cancer diagnosis. Controls were matched on age, gender, time to censoring (death or end of follow-up), and source of imaging referral (emergency, inpatient or outpatient departments). Images were obtained using 3 types of scanners (GE Picker PQ 6000, Siemens AS plus, and Toshiba Aquilion 16) using standard protocols. The chest CTs were obtained in the late arterial phase, following a 30 second

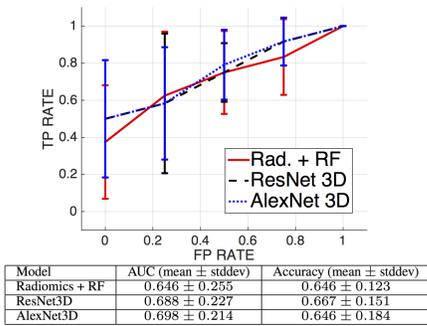
delay after the administration of intravenous contrast (Omnipaque350/Ultravist370), and were annotated by a radiologist using semi-automated segmentation tools contained in the Vitrea software suite (Vital Images, Toshiba), where the following anatomies have been segmented: muscle, body fat, aorta, spinal column, epicardial fat, heart, and lungs.

The evaluation of the methodologies is based on a 6-fold cross-validation experiment, where each fold contains 20 cases and 20 matched controls for training and 4 cases and 4 matched controls for testing. The classification performance is measured using the mean accuracy over the six experiments, with accuracy computed by  $\frac{TP+TN}{TP+FP+TN+FN}$ , where  $TP$  represents correct mortality prediction,  $TN$  denotes correct survival prediction,  $FP$  means incorrect mortality prediction, and  $FN$ , incorrect survival prediction. We also show the receiver operating characteristic (ROC) curve and area under curve (AUC) [17] using the classifier confidence on the 5-year mortality classification.

For the radiomics method, we hand-crafted 16210 features, where 2506 features come from the aorta, 2506 from heart, 2236 from lungs, 2182 from epicardial fat, 2182 from body fat, 2182 from muscle, and 2416 from spinal column<sup>2</sup>, where 936 represent domain knowledge features [6, 7, 15] (see Sec. 3). For classification, we used random forests (RF) [18], trained with 900 trees, minimum *nodesize* of 5 (minimum number of training samples per node), and with *mtry* of 3 (i.e., number of variables sampled as candidates for each node split). Note that in our previous work [19], we showed empirically that this configuration produced the best results for the radiomics method.

We test two types of ConvNets, extended to 3D inputs: AlexNet3D [9] and ResNet3D [10]. AlexNet3D [9] has four convolutional layers, where the input has eight 3-D channels (chest CT and 7 segmentation maps), the first layer has 50 filters and the second to fourth layers have 100 filters of size  $5 \times 5 \times 2$  (i.e., these are 3-D filters). The first convolutional layer has ReLU activation [20], the fifth layer contains 6000 nodes, and the output layer has two nodes. For training, dropout [21] of 0.35 is applied to all layers, the learning rate starts at 0.0005, from epochs 1 to 10, which is then continuously reduced until it reaches 0.00001 from epochs 60 to 120, and we use RMS prop [22] with  $\rho = 0.9$ , and  $\epsilon = 10^{-6}$ . ResNet3D [10] uses the same eight 3-D input channels as AlexNet3D, where the first convolutional layer has  $32 \times 3 \times 3$  filters with ReLU activation [20], followed by a max-pooling layer that reduces the volume by a factor of two. Then we have three residual blocks (first stage composed of a convolutional layer with  $32 \times 3 \times 3$  filters, followed by two convolutional layers with  $32 \times 3 \times 3$  filters, whose output is summed to the output of the first stage of this block) with 32 filters, then four residual blocks with 64 filters, then six residual blocks with 128 filters, then three residual blocks with 256 filters, and the output layer has two nodes. For training, we apply dropout [21] of 0.20 at the beginning of each residual block, and use the same configuration as in AlexNet3D in terms of the learning rate and RMS prop [22]. These models are implemented on Theano + Lasagne [23].

<sup>2</sup>Most of these features are hand-crafted with the methodology provided by J. Carlson (<https://cran.r-project.org/web/packages/radiomics/>).



**Fig. 2.** Mean/standard deviation of the ROC (graph on top), AUC and accuracy (table at the bottom) of the experiments on the testing set using deep learning and radiomics methods (results are slightly different from [19] due to the use of different cross validation sets).

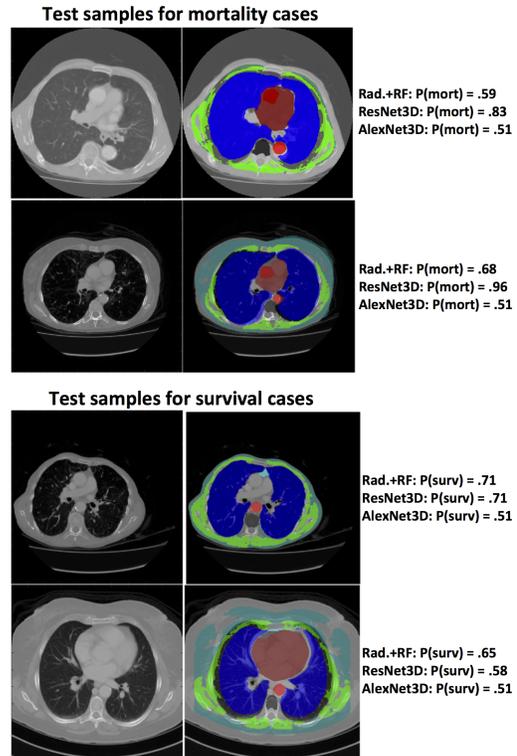
**Results:** We show the mean and the standard deviation of the ROC curves for the testing set of the radiomics (+ RF), ResNet3D, and AlexNet3D models in Fig. 2, which also shows a table with the mean and standard deviation of the AUC and accuracy of the testing set of the deep learning and the radiomics models<sup>3</sup>. Using the t-test for paired samples, we note that there is no significant difference between any pair of models in terms of accuracy and AUC results on the testing set. Finally, in Fig. 3, we show four chest CT examples with the output from both models.

## 5. DISCUSSION AND CONCLUSIONS

The experiments demonstrate promising results, with prediction accuracy from routinely obtained chest CTs similar to the current state-of-the-art clinical risk scores, despite our small dataset and our exclusion of highly predictive covariates such as age and gender. Furthermore, expert review of the correctly classified images (such as the example cases in Fig. 3) suggests that our models may be identifying medically plausible imaging biomarkers. The comparison between deep learning and radiomics models shows that they produce comparable classification results, but the deep learning model offers several advantages, such as automatic feature learning, and unified feature and classifier learning.

These advantages mitigate the issues of hand-crafting features, which requires expert domain knowledge, and the complicated multi-stage learning process of radiomics. It is remarkable that a deep learning model implemented with relative simplicity could produce competitive results compared to the radiomics method, which uses features that have been heavily tuned for the task at hand [6, 7, 15], and relies on an extensive set of initial features (e.g., we have 16210 features). This hand-crafting task would need to be re-tuned for every new problem in radiomics, unlike the deep learning approach. A possible question that one might have is the need of the segmentation maps. We actually tested that hypothesis (but results are omitted due to lack of space) and results confirmed

<sup>3</sup>While AlexNet 3D did not present overfitting issues, ResNet 3D and radiomics + RF presented varying degrees of overfitting issues that need to be investigated further.



**Fig. 3.** Testing examples of 5-year mortality classification (see on the right, P(mort) and P(surv) results for the probability of mortality and survival) produced by the Radiomics (+RF), ResNet3D and AlexNet3D, where the image on the left shows a mid-level plane of the chest CT, and the image on the right displays the same plane overlaid with the following annotations: aorta (red), spinal column (gray), epicardial fat (cyan), fat (light cyan), heart (light red), lungs (blue), muscle (green). The mortality examples in general show significant coronary artery calcification, enlarged aortic root and heart, low bone density, and muscle mass loss.

that the use of segmentation maps produce significantly better classification results. Finally, we believe that the deep learning results can be improved with the use of pre-training [8, 9] and both models would benefit significantly from the integration of predictive epidemiological information (e.g., gender and age).

In this paper, we show the first proof of concept experiments for a system capable of predicting 5-year mortality in elderly individuals from chest CTs alone. The widespread use of medical imaging suggests that our methods will be clinically useful after being successfully tested in large scale problems (in fact, we are in the process of acquiring larger annotated datasets), as the only required inputs are already highly utilised: the medical images. We also note that the proposed deep learning model can be easily extended to other important medical outcomes, and other imaging modalities.

## 6. REFERENCES

- [1] Andrea Ganna and Erik Ingelsson, “5 year mortality predictors in 498 103 uk biobank participants: a prospective population-based study,” *The Lancet*, vol. 386, no. 9993, pp. 533–540, 2015.
- [2] Lindsey C Yourman, Sei J Lee, et al., “Prognostic indices for older adults: a systematic review,” *Jama*, vol. 307, no. 2, pp. 182–192, 2012.
- [3] Hugo JWL Aerts, Emmanuel Rios et al., “Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach,” *Nature communications*, vol. 5, 2014.
- [4] Philippe Lambin, Emmanuel Rios-Velazquez, et al., “Radiomics: extracting more information from medical images using advanced feature analysis,” *European Journal of Cancer*, vol. 48, no. 4, pp. 441–446, 2012.
- [5] Virendra Kumar, Yuhua Gu, et al., “Radiomics: the process and the challenges,” *Magnetic resonance imaging*, vol. 30, no. 9, pp. 1234–1248, 2012.
- [6] Jan S Bauer, Tobias D Henning, et al., “Volumetric quantitative ct of the spine and hip derived from contrast-enhanced mdct: conversion factors,” *American Journal of Roentgenology*, vol. 188, no. 5, pp. 1294–1301, 2007.
- [7] Akane Haruna, Shigeo Muro, et al., “Ct scan findings of emphysema predict mortality in copd,” *CHEST Journal*, vol. 138, no. 3, pp. 635–640, 2010.
- [8] Geoffrey E Hinton and Ruslan R Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [9] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” *arXiv preprint arXiv:1512.03385*, 2015.
- [11] Yaniv Bar, Idit Diamant, Lior Wolf, and Hayit Greenspan, “Deep learning with non-medical training used for chest pathology identification,” in *SPIE Medical Imaging*. International Society for Optics and Photonics, 2015, pp. 94140V–94140V.
- [12] Dan Ciorean, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber, “Deep neural networks segment neuronal membranes in electron microscopy images,” in *Advances in neural information processing systems*, 2012, pp. 2843–2851.
- [13] Neeraj Dhungel, Gustavo Carneiro, and Andrew P Bradley, “Deep learning and structured prediction for the segmentation of mass in mammograms,” in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, pp. 605–612. Springer, 2015.
- [14] Gustavo Carneiro, Jacinto Nascimento, and Andrew P Bradley, “Unregistered multiview mammogram analysis with pre-trained deep learning models,” in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, pp. 652–660. Springer, 2015.
- [15] Khurram Nasir, Jonathan Rubin, et al., “Interplay of coronary artery calcification and traditional risk factors for the prediction of all-cause mortality in asymptomatic individuals,” *Circulation: Cardiovascular Imaging*, vol. 5, no. 4, pp. 467–473, 2012.
- [16] Yann LeCun and Yoshua Bengio, “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, 1995.
- [17] Trevor Hastie, Robert Tibshirani, Jerome Friedman, and James Franklin, “The elements of statistical learning: data mining, inference and prediction,” *The Mathematical Intelligencer*, vol. 27, no. 2, pp. 83–85, 2005.
- [18] Leo Breiman, “Random forests,” *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [19] Gustavo Carneiro, Luke Oakden-Rayner, Andrew P Bradley, Jacinto Nascimento, and Lyle Palmer, “Automated 5-year mortality prediction using deep learning and radiomics features from chest computed tomography,” *arXiv preprint arXiv:1607.00267*, 2016.
- [20] Vinod Nair and Geoffrey E Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 807–814.
- [21] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [22] Yann N Dauphin, Harm de Vries, Junyoung Chung, and Yoshua Bengio, “Rmsprop and equilibrated adaptive learning rates for non-convex optimization,” *arXiv preprint arXiv:1502.04390*, 2015.
- [23] The Theano Development Team, Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, et al., “Theano: A python framework for fast computation of mathematical expressions,” *arXiv preprint arXiv:1605.02688*, 2016.