



## King's Research Portal

DOI:

[10.1007/978-3-030-39074-7\\_3](https://doi.org/10.1007/978-3-030-39074-7_3)

*Document Version*

Peer reviewed version

[Link to publication record in King's Research Portal](#)

*Citation for published version (APA):*

Puyol-Antón, E., Ruijsink, B., Clough, J. R., Oksuz, I., Rueckert, D., Razavi, R., & King, A. P. (2020). Assessing the Impact of Blood Pressure on Cardiac Function Using Interpretable Biomarkers and Variational Autoencoders. In M. Pop, M. Sermesant, O. Camara, X. Zhuang, S. Li, A. Young, T. Mansi, & A. Suinesiaputra (Eds.), *Statistical Atlases and Computational Models of the Heart. Multi-Sequence CMR Segmentation, CRT-EPIggy and LV Full Quantification Challenges - 10th International Workshop, STACOM 2019, Held in Conjunction with MICCAI 2019, Revised Selected Papers* (pp. 22-30). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 12009 LNCS). SPRINGER. [https://doi.org/10.1007/978-3-030-39074-7\\_3](https://doi.org/10.1007/978-3-030-39074-7_3)

### **Citing this paper**

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

### **General rights**

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

### **Take down policy**

If you believe that this document breaches copyright please contact [librarypure@kcl.ac.uk](mailto:librarypure@kcl.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

# Assessing the Impact of Blood Pressure on Cardiac Function Using Interpretable Biomarkers and Variational Autoencoders

Esther Puyol-Antón<sup>\*1</sup>, Bram Ruijsink <sup>\*1,2</sup>, James R. Clough<sup>1</sup>, Ilkay Oksuz<sup>1</sup>, Daniel Rueckert<sup>3</sup>, Reza Razavi<sup>1,2</sup>, and Andrew P. King<sup>1</sup>

<sup>1</sup> School of Biomedical Engineering & Imaging Sciences, King's College London, UK.

<sup>2</sup> St Thomas' Hospital NHS Foundation Trust, London, UK.

<sup>3</sup> Biomedical Image Analysis Group, Imperial College London, UK.

**Abstract.** Maintaining good cardiac function for as long as possible is a major concern for healthcare systems worldwide and there is much interest in learning more about the impact of different risk factors on cardiac health. The aim of this study is to analyze the impact of systolic blood pressure (SBP) on cardiac function while preserving the interpretability of the model using known clinical biomarkers in a large cohort of the UK Biobank population. We propose a novel framework that combines deep learning based estimation of interpretable clinical biomarkers from cardiac cine MR data with a variational autoencoder (VAE). The VAE architecture integrates a regression loss in the latent space, which enables the progression of cardiac health with SBP to be learnt. Results on 3,600 subjects from the UK Biobank show that the proposed model allows us to gain important insight into the deterioration of cardiac function with increasing SBP, identify key interpretable factors involved in this process, and lastly exploit the model to understand patterns of positive and adverse adaptation of cardiac function.

**Key words:** Cardiac function; variational autoencoder; cardiac risk factors

## 1 Introduction

Preventing the development of heart disease in patients with known risk factors, such as hypertension, represents a major challenge for healthcare systems worldwide. Although much is known about how these risk factors influence development of disease, the recent availability of large scale databases such as the UK Biobank represents an excellent opportunity to extend this knowledge. Learning from a large number of highly detailed, multidimensional cardiac magnetic resonance (CMR) datasets can help further understanding of how risk factors impact cardiac function, and tailor medical interventions to individual patients. Traditionally, machine learning techniques have relied on the use of handcrafted features to effectively perform a specific task without using explicit instructions.

---

\* Joint first authors.

In some cases, the accuracy of these models was limited by the model being restricted to the use of these features. Recently, deep learning (DL) techniques have demonstrated a significant increase in performance over traditional machine learning methods. DL allows features to be learned from the data themselves, without preselection. One drawback of DL approaches is the lack of interpretability, as the learned relationships and features are often abstract and opaque to human users. Especially in medicine, interpretability and accountability are vital for two main reasons: (a) they can promote clinician trust in the learned models; and (b) the use of well-established, interpretable biomarkers allows the models to be used to better understand disease processes, translate the results to other populations and to interpret the newly learned information in the light of already existing clinical scientific evidence. For example, complex full 3D cardiac motion of the heart can be used to outperform current models in survival estimation for patients with pulmonary hypertension [1]. While methods like this demonstrate the power of DL, it is difficult for clinicians to understand the features underlying the predictions and to use the model to better understand the disease.

A large number of biomarkers can be calculated from CMR. These are well-understood by clinicians and provide comprehensive information about underlying physiological processes. However, estimating them all is labour-intensive. In this paper, we employ a fully automated DL-based pipeline for estimating a wide range of biomarkers of cardiac function from cine CMR data. Our main contribution is to propose a framework that enables the interpretability of these automatically computed biomarkers to be combined with the power of learned features in DL. This framework is based on the use of a variational autoencoder with a latent space regression loss (R-VAE), in which the input data are the clinical biomarkers. In addition, we use a dummy variable in the regression to differentiate between population groups. We use the proposed method to investigate the impact of systolic blood pressure (SBP: a measure of hypertension) on cardiac function in the healthy population differentiated by gender.

**Related Work:** In the clinical literature, many groups have investigated the relationship between SBP and ventricular structure, function and geometry [2, 3]. However, most studies only investigated the influence of SBP on global parameters or only in the left ventricle. The proposed pipeline enables a more detailed investigation of the impact of SBP on cardiac function, as we demonstrate in Section 4. VAEs have previously been used for identification and visualization of features in medical image-based classification tasks [4–6], but the features were still learnt from the data and were not well-established clinical biomarkers as in our work. DL models have also previously been proposed for regression using interpretable features. For example, Xie *et al.* [7] proposed an autoencoder-based deep belief regression network to forecast daily particulate matter concentrations. Similarly, Bose *et al.* [8] proposed a stacked autoencoder based regression framework to optimize process control and productivity in intelligent manufacturing. Both techniques combined handcrafted features obtained from the image domain with autoencoders. In the medical field, Xie *et al.* [9] have proposed a deep autoencoder model for regression of gene expression profiles from genotype.


Similar to these works we integrate a regression loss into the autoencoder to learn relationships between the latent space and another variable (SBP in our case). Our work is methodologically distinct from [7–9] as we employ a VAE, which enables us to sample from the distribution of the latent space and decode the clinical biomarkers for these samples. We also employ a dummy variable in the regression to enable the investigation to be stratified by gender. Our work also represents the first time that a regression-based autoencoder has been applied to investigate the impact of risk factors on cardiac function.

## 2 Materials

We evaluate our approach on subjects selected from the UK Biobank data set, which contains multiple imaging and non-imaging information from more than half a million 40-69 year-olds. We included only participants with CMR imaging data. From this group, we excluded participants with a history of cardiovascular disease, respiratory disease, renal disease, cancer, rheumatological disease, symptoms of chest pain or dyspnoea. For each subject, the following cine CMR acquisitions were used; a short-axis (SA) stack covering the full heart, and two orthogonal long-axis (LA) acquisitions (2-chamber (2Ch) and 4-chamber (4Ch) views). For each image slice 50 temporal frames were available covering a full cardiac cycle (temporal resolution  $\approx 14\text{-}24$  ms/frame). All CMR imaging was carried out on a 1.5 Tesla scanner (Siemens Healthcare, Erlangen, Germany). Details of the image acquisition protocol can be found in [10]. Blood pressure was measured using the HEM-70151T digital blood pressure monitor (Omron, Hoofddorp, The Netherlands) [11]

## 3 Methods

In the following sections we first describe the automated estimation of biomarkers of cardiac function from CMR images, and secondly present the R-VAE network used to learn the relationship between cardiac function and SBP. Figure 1 summarizes these steps and how they interrelate.



**Fig. 1.** Overview of the proposed framework for a VAE regression model based on automatically estimated clinical biomarkers.

### 3.1 Biomarker Estimation From CMR

The procedure used to extract the clinical biomarkers of cardiac function is based on the work published in [12] and is briefly summarized below.

**Automatic segmentation network:** We first used a fully-convolutional network with a 17 convolutional layer VGG-like architecture for automatic segmentation of the left ventricle (LV) blood pool, LV myocardium and right ventricle (RV) blood pool from SA and LA slices in all frames through the cardiac cycle [13, 14]. After this, all segmentations were aligned to correct for breath-hold induced motion artefacts using the iterative registration algorithm proposed in [14].

**Biomarker calculation:** LV and RV blood volume curves were calculated from the obtained segmentations. From these curves, end-diastolic volume (EDV), end-systolic volume (ESV), stroke volume (SV), ejection fraction (EF), LVED mass (LVEDM), peak ejection rate (PER), peak early filling rate (PEFR), atrial contribution (AC) and peak atrial filling rate (PAFR) were obtained. Cardiac volumes were indexed to body surface area (BSA) using the Dubois and Dubois formula [15]. The complete list of biomarkers calculated was: iLVEDV (indexed LVEDV), iLVSV, LVEDM, LVPER, LVPFR, LVPAFR, LVAC, iRVEDV, RVPER, RVPFR, RVPAFR and RVAC.

**Quality control:** Similar to [12], two quality control (QC) methods were implemented to automatically reject subjects with insufficient image quality or incorrect segmentations, ensuring the robustness of the estimated biomarkers. The first QC step (QC1) used trained DL models to reject any image with poor quality or incorrect planning, and the second QC step (QC2) detected incorrect segmentation results using an SVM model that identified physiologically unrealistic or unusual volume curves.

### 3.2 Deep Learning Regression Model

To combine the interpretability of handcrafted features with the power of DL we propose to use a VAE featuring a regression loss in the latent space to simultaneously learn efficient representations of cardiac function and map their change with regard to differences in SBP. As a regression model we used the multivariable regression modelling commonly used in epidemiological studies [16], where the effect of different independent variables are included as confounders on the regression model. As a result, the VAE linearizes the relationship between different clinical biomarkers and the variable of interest (SBP in this case), and these features are incorporated in a standard regression model.

**Variational autoencoder:** The encoding part of a VAE allows a number of features,  $\mathbf{x}$ , to be mapped into a lower dimensional representation (the *latent space*), whilst the decoder maps this representation back to the original higher dimensional space. The proposed VAE has  $N = 13$  input units representing the clinical biomarkers, two hidden layers with 8 and 4 hidden units respectively and a latent space of dimensionality 2. To avoid over-fitting we apply dropout with probability 0.3 after each hidden layer during training.

**Linear regression with indicator (dummy) variable:** We use a linear regression model in the latent space to estimate SBP, that incorporates a dummy variable (encoded as 0 or 1) to differentiate between population groups. In our experiments we used gender as a dummy variable, but the method is general and could be used to investigate a wide range of other factors. Mathematically, the linear regression model can be formulated as follows:  $y = w^T(z + D) + \epsilon$ , where  $z$  are the latent space activations,  $D$  the dummy variable,  $y$  the ground truth label (i.e. SBP) and  $w$  the regression coefficients. The regression loss is the

$$\text{mean squared error: } L_{\text{regression}} = \frac{1}{N} \sum_{i=1}^N (y_i - w_i^T(z_i + D_i))^2.$$

**Joint learning for regression:** We denote the input data by  $\mathbf{x} = [x_1, x_2, \dots, x_N]$  (i.e. a vector of  $N = 13$  clinical biomarkers) and its corresponding latent space representation as  $\mathbf{z} = [z_1, z_2]$ . The decoded clinical biomarkers are denoted by  $\hat{\mathbf{x}} = [\hat{x}_1, \hat{x}_2, \dots, \hat{x}_N]$ , and the predicted label by  $\hat{y} = \text{Regressor}(\mathbf{z})$ . We combine the VAE loss and the regression loss by minimising the following joint loss function:

$$L_{\text{R-VAE}} = L_{\text{recon}} + \alpha L_{\text{KL}} + \beta L_{\text{regression}} \quad (1)$$

where  $\alpha$  and  $\beta$  control the weights of the components of the loss function,  $L_{\text{KL}}$  is the Kullback-Leibler divergence between the learnt latent distribution and a unit Gaussian, and  $L_{\text{regression}}$  is the Huber loss for the regression task. We first train the model only using the VAE loss, i.e.  $\beta = 0$ , and secondly train both the VAE and the regression together using  $\beta = 2$ . We set  $\alpha = 0.3$  throughout.

## 4 Experiments and Results

Using our selection criteria, 3,781 subjects were included in our experiments. Of these, 54 subjects were rejected during QC1, and a further 127 during QC2. The remaining 3,600 cases were used to build the models. Of these cases, 1,321 had normotension (SBP <120 mmHg), 1,697 hypertension (SBP >140 mmHg) and 582 cases had prehypertension (a SBP between 120 and 140 mmHg).

### Experiment 1 - Comparative evaluation on regression task:


We compared our proposed R-VAE model with two state-of-the-art techniques for multivariate regression: (1) Lasso regression, a linear regression model with a  $l_1$  regularizer [17]; and (2) Random Forest regression [18], an ensemble method that has shown excellent performance in complex regression and classification tasks. For all methods, we used a five-fold cross validation to obtain the optimal model. We split the dataset into training, validation and test (60/20/20), and optimized the hyperparameters using a grid search strategy. We calculated the root-mean-square deviation (RMSD), the normalized root-mean-square deviation (nRMSD), and the coefficient of determination ( $R^2$ ) between the ground truth SBP and the predicted SBP. Table 1 shows these results. It can be seen that all methods performed similarly with regard to regression. However, note that the next two experiments are only made possible by our use of the R-VAE architecture and would not be possible with the other regression techniques.

**Table 1.** Comparison of R-VAE, Lasso regression and Random Forest regression.

Methods	RMSD (mmHg)	nRMSD (mmHg)	$R^2$ (%)
Lasso [17]	13.2	0.11	0.35
Random Forest [18]	12.98	0.12	0.33
R-VAE	11.36	0.10	0.69


**Experiment 2 - Investigating the effect of SBP on cardiac function:**

We applied our R-VAE network to investigate how cardiac function changes with increasing SBP in healthy individuals, stratified by gender. We applied our R-VAE network to investigate how cardiac function changes with increasing SBP in healthy individuals, stratified by gender. We sampled the latent space of the model along the regression line for different SBP values (between 100 to 170 mmHg in steps of 10 mmHg). At each step, we took 20 random samples from a normal distribution in a perpendicular direction to the regression line and used the R-VAE to decode the clinical biomarkers. Figure 2 shows the means and standard deviations for a selection of the decoded biomarkers stratified by gender. The results show that iLVEDV (indexed LVEDV) decreases with increasing SBP, while iRVEDV (indexed RVEDV) remains constant. LVPAFR and LVPER increase with increasing SBP. For both iLVEDV and LVPER, the change seems to be larger in males compared to females. Overall, the observed changes in the models' predictions suggest that parameters associated with diastolic function of the LV are mostly affected by SBP, while the RV was less affected. These results suggest that stiffening of the LV myocardium could be an important disease process in deterioration of cardiac function in the light of increased SBP.

**Fig. 2.** SBP-related changes in iLVEDV, iRVEDV, LVPAFR and LVPER. Red represents females and blue represents males. Bars represent standard deviations. Black dotted lines represent the linear tendency curves between the cardiac biomarkers and ground-truth SBP data.**Experiment 3 - Identifying abnormal response:**

In the normal population, some individuals with prehypertension might be predisposed to increased risk of cardiac disease, while others are not. We used our R-VAE model to identify subjects from the prehypertension group in whom predicted SBP was lower (i.e. predicted normotension) or higher (i.e. predicted hypertension) based on the latent space regression, assuming that being wrongly

classified as normotensive or hypertensive identifies individuals with low versus increased risk of developing cardiac disease. Subsequently, we decoded the cardiac biomarkers for these cases using latent features at the regression line of their true SBP as well as using latent features at their predicted SBP. We calculated the percentage difference for each biomarker for the cases of under- and over-prediction respectively and investigated which factors contributed most to the lower or higher prediction in these subjects. Figure 3 shows the mean difference of selected biomarkers that lead to classification as normotensive (dark) and hypertensive (light). Biomarkers related to LV diastolic function (blue bars) showed the largest changes with regard to the under- and overprediction. These results show again that diastolic function of the LV was a major contributor to the model predictions of SBP. Moreover, they suggest that biomarkers related to LV diastolic function might be effective when trying to stratify cardiac risk, in particular in subjects with ‘prehypertension’.



**Fig. 3.** Mean change (percentage) of each biomarker in prehypertension cases that were classified by the regression model as normotensive (dark) and hypertensive (light) with respect to values predicted by the model using the actual observed SBP. Values further away from zero mean a larger impact of these biomarkers.

## 5 Discussion

In this paper, we have presented an automated DL method for analysing cardiac function and predicting cardiac risk profiles from CMR. Our framework encompasses all steps from CMR segmentation through quality control to modelling of



the impact of SBP, a known cardiovascular risk factor, on cardiac function. Instead of inputting CMR images directly into the R-VAE, we chose to first automatically estimate clinical biomarkers from the images. Combined, these biomarkers give a comprehensive description of cardiac function and are well-understood by clinicians. While some of the information of the high dimensional image data is inevitably lost by this approach, it allows the model to be interpretable by clinicians directly.

The combination of the VAE with the regression loss allowed us to decode the clinical biomarkers from the latent space, while also providing a mapping of these biomarkers to another variable, SBP. As we show in Experiment 2, this design enabled us to get a clear description of the changes in cardiac function that occur with increasing SBP. Using the trained model, we showed that increases in SBP are mainly linked to changes in diastolic LV function. This suggests that SBP results in slowly progressive changes in the myocardium that increase ventricular stiffening. As shown in Experiment 3, the model also allowed us to identify key factors separating high and low risk subjects. Again, due to the interpretability of the framework, this allowed us to identify biomarkers that could be further investigated for their utility in screening patients in clinical practice. SBP is not the only factor influencing cardiac function and that explains the relatively low  $R^2$  of the regression models. In this paper, we used SBP as an example to illustrate the potential power of our proposed method. We aim to further extend our model in the following ways: we plan to include image and segmentation data directly into the model, in combination with the clinical biomarkers to maintain interpretability; we also plan to extend the model to investigate more risk factors. In conclusion, this work represents a novel use of DL which has produced an important contribution to furthering our understanding of the influences on cardiac function.

## Acknowledgements

This work was supported by the EPSRC (grants EP/R005516/1 and EP/P001009/1) and the Wellcome EPSRC Centre for Medical Engineering at the School of Biomedical Engineering and Imaging Sciences, Kings College London (WT 203148/Z/16/Z). This research has been conducted using the UK Biobank Resource under Application Number 17806.

## References

1. Bello et al., Deep-learning cardiac motion analysis for human survival prediction. *Nat Mach Intell*, 1(2):95, 2019.
2. Bajpai et al., Impact of prehypertension on left ventricular structure, function and geometry. *J Clin Diagn Res* 8(4):BC07, 2014.
3. Mo et al., The Bergen Blood Pressure Study: prehypertensive changes in cardiac structure and function in offspring of hypertensive families. *Blood pressure* 4(1):16-22,1995.

4. Biffi et al. Learning interpretable anatomical features through deep generative models: Application to cardiac remodeling. MICCAI 464-471, 2018.
5. Biffi et al. Explainable Shape Analysis through Deep Hierarchical Generative Models: Application to Cardiac Remodeling, 2019. (arXiv preprint arXiv:1907.00058).
6. Clough et al. Global and Local Interpretability for Cardiac MRI Classification. MICCAI, 2019 in press (arXiv preprint arXiv:1906.06188).
7. Xie et al., Autoencoder-based deep belief regression network for air particulate matter concentration forecasting, *J Intel Fuzzy Syst*, 34(6):3475-86, 2018.
8. Bose et al., Machine Load Estimation Via Stacked Autoencoder Regression, ICASSP 2126-2130, 2018.
9. Xie et al., A deep auto-encoder model for gene expression prediction. *BMC genomics* 18(9), 2017
10. Petersen et al., UK biobank's cardiovascular magnetic resonance protocol, *JCMR*, 18(1): 1-8, 2016.
11. Chan et al. Associations with intraocular pressure in a large cohort: results from the UK Biobank. *Ophthalmology* 123.4:771-782, 2016.
12. Ruijsink and Puyol-Antón et al., Fully automated, quality-controlled ventricular assessment from CMR: validation and large-scale application to characterize cardiac function, *JACC: Cardiovascular Imaging*, 2019
13. Bai et al., Automated cardiovascular magnetic resonance image analysis with fully convolutional networks, *JCMR* 20(65), 2018.
14. Sinclair et al., Fully Automated segmentation-based respiratory motion correction of multiplanar cardiac magnetic resonance images for large-scale datasets, MICCAI 332-340, 2017.
15. Du Bois D., Du Bois EF. A formula to estimate the approximate surface area if height and weight be known. *Arch Intern Med* 17(6):86371, 1916.
16. Thompson, S.G. and Higgins, J.P. How should metaregression analyses be undertaken and interpreted?. *Statistics in medicine* 21.11:1559-1573, 2002.
17. Tibshirani et al., Regression shrinkage and selection via the lasso, *J R Stat Soc Series B Stat Methodol* 58(1), 1996.
18. Lin et al., Random forests and adaptive nearest neighbors, *J Am Stat Assoc* 101(474), 2006.