

Published in final edited form as:

Nat Mach Intell. 2019 February 11; 1: 95–104. doi:10.1038/s42256-019-0019-2.

Deep learning cardiac motion analysis for human survival prediction

Ghalib A. Bello^{#1}, Timothy J.W. Dawes^{#1,2}, Jinming Duan^{1,3}, Carlo Biffi^{1,3}, Antonio de Marvao¹, Luke S. G. E. Howard⁴, J. Simon R. Gibbs^{2,4}, Martin R. Wilkins⁶, Stuart A. Cook^{1,2,5}, Daniel Rueckert³, and Declan P. O'Regan^{1,*}

¹MRC London Institute of Medical Sciences, Imperial College London, UK

²National Heart and Lung Institute, Imperial College London, UK

³Department of Computing, Imperial College London, UK

⁴Imperial College Healthcare NHS Trust, London, UK

⁵National Heart Centre Singapore, Singapore, and Duke-NUS Graduate Medical School, Singapore

⁶Division of Experimental Medicine, Department of Medicine, Imperial College London, UK

These authors contributed equally to this work.

Abstract

Motion analysis is used in computer vision to understand the behaviour of moving objects in sequences of images. Optimising the interpretation of dynamic biological systems requires accurate and precise motion tracking as well as efficient representations of high-dimensional motion trajectories so that these can be used for prediction tasks. Here we use image sequences of the heart, acquired using cardiac magnetic resonance imaging, to create time-resolved three-dimensional segmentations using a fully convolutional network trained on anatomical shape priors. This dense motion model formed the input to a supervised denoising autoencoder (4D*survival*), which is a hybrid network consisting of an autoencoder that learns a task-specific latent code representation trained on observed outcome data, yielding a latent representation optimised for survival prediction. To handle right-censored survival outcomes, our network used a Cox partial likelihood loss function. In a study of 302 patients the predictive accuracy (quantified by Harrell's C-index) was significantly higher ($p = .0012$) for our model $C=0.75$ (95% CI: 0.70 - 0.79) than the human benchmark of $C=0.59$ (95% CI: 0.53 - 0.65). This work demonstrates how a complex

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

* declan.oregan@imperial.ac.uk.

Author contributions statement (CRediT)

G.A.B., C.B. and T.J.W.D. - methodology, software, formal analysis and writing original draft. J.D. - methodology, software and writing original draft; A. de M. - formal analysis; L.S.G.E.H., J.S.R.G, M.R.W and S.A.C. - investigation; D.R. - software and supervision; D.P.O'R. - conceptualization, supervision, writing – review and editing, funding acquisition. All authors reviewed the final manuscript.

Competing interests

The authors declare no competing financial interests.

computer vision task using high-dimensional medical image data can efficiently predict human survival.

Keywords

Survival; Machine Learning; Magnetic Resonance Imaging, Cine; Hypertension, Pulmonary; Heart Failure; Motion

Introduction

Techniques for vision-based motion analysis aim to understand the behaviour of moving objects in image sequences.¹ In this domain deep learning architectures have achieved a wide range of competencies for object tracking, action recognition, and semantic segmentation.² Making predictions about future events from the current state of a moving three dimensional (3D) scene depends on learning correspondences between patterns of motion and subsequent outcomes. Such relationships are important in biological systems which exhibit complex spatio-temporal behaviour in response to stimuli or as a consequence of disease processes. Here we use recent advances in machine learning for visual processing tasks to develop a generalisable approach for modelling time-to-event outcomes from time-resolved 3D sensory input. We tested this on the challenging task of predicting survival due to heart disease through analysis of cardiac imaging.

The motion dynamics of the beating heart are a complex rhythmic pattern of non-linear trajectories regulated by molecular, electrical and biophysical processes.³ Heart failure is a disturbance of this coordinated activity characterised by adaptations in cardiac geometry and motion that lead to impaired organ perfusion.⁴ For this prediction task we studied patients diagnosed with pulmonary hypertension (PH), characterised by right ventricular (RV) dysfunction, as this is a disease with high mortality where the choice of treatment depends on individual risk stratification.⁵ Our input data were derived from cardiac magnetic resonance (CMR) which acquires imaging of the heart in any anatomical plane for dynamic assessment of function. While explicit measurements of performance obtained from myocardial motion tracking detect early contractile dysfunction and act as discriminators of different pathologies,^{6, 7} we hypothesized that learned features of complex 3D cardiac motion would provide enhanced prognostic accuracy.

A major challenge for medical image analysis has been to automatically derive quantitative and clinically-relevant information in patients with disease phenotypes. Our method employs a fully convolutional network (FCN) to learn a cardiac segmentation task from manually-labelled priors. The outputs are smooth 3D renderings of frame-wise cardiac motion which are used as input data to a supervised denoising autoencoder prediction network which we refer to as *4Dsurvival*. The aim is to learn latent representations robust to noise and salient for survival prediction. We then compared our model to a benchmark of conventional human-derived volumetric indices and clinical risk factors in survival prediction.

Results

Baseline Characteristics

Data from all 302 patients with incident PH were included for analysis. Objective diagnosis was made according to haemodynamic and imaging criteria.⁵ Patients were investigated between 2004 and 2017, and were followed-up until November 27, 2017 (median 371 days). All-cause mortality was 28% (85 of 302). Table 1 summarizes characteristics of the study sample at the date of diagnosis. No subjects' data were excluded.

MR Image Processing

Automatic segmentation of the ventricles from gated CMR images was performed for each slice position at each of 20 temporal phases producing a total of 69,820 label maps for the cohort (Figure 1A). Image registration was used to track the motion of corresponding anatomic points. Data for each subject was aligned producing a dense model of cardiac motion across the patient population (Figure 1B) which was then used as an input to the *4Dsurvival* network.

Predictive performance

Bootstrapped internal validation was applied to the *4Dsurvival* and benchmark models. The apparent predictive accuracy for *4Dsurvival* was $C = 0.86$ and the optimism-corrected value was $C = 0.75$ (95% CI: 0.70-0.79). The *4Dsurvival* model out-performed benchmark models of volumetric CMR parameters ($p = 0.0012$): apparent predictive accuracy $C = 0.60$ and optimism-adjusted $C = 0.59$ (95% CI: 0.53-0.65); myocardial strain parameters ($p = 0.016$): apparent predictive accuracy $C = 0.64$ and optimism-adjusted $C = 0.61$ (95% CI: 0.57-0.66); and a joint analysis of both imaging and clinical risk factors ($p = 0.006$): apparent predictive accuracy $C = 0.66$ and optimism-adjusted $C = 0.64$ (95% CI: 0.57-0.70). Figure 2 shows Kaplan-Meier plots which depict the survival probability estimates over time, stratified by risk groups defined by each model's predictions (see Supplementary section for details). After bootstrap validation, a final model was created using the training and optimization procedure outlined in the Methods section (optimal hyperparameters for this model are summarized in Table 2).

Visualization of Learned Representations

To assess the ability of the *4Dsurvival* network to learn discriminative features from the data, we examined the encoded representations by projection to 2D space using Laplacian Eigenmaps⁸ (Figure 3A). In this figure, each subject is represented by a point, the colour of which is based on the subject's survival time, i.e. time elapsed from baseline (date of MRI scan) to death (for uncensored patients), or to the most recent follow-up date (for censored patients). Survival time was truncated at 7 years for ease of visualization. As is evident from the plot, our network's compressed representations of 3D motion input data show distinct patterns of clustering according to survival time. Figure 3A also shows visualizations of RV motion for 2 exemplar subjects at opposite ends of the risk spectrum. We also assessed the extent to which motion in various regions of the RV contributed to overall survival prediction. Fitting univariate linear models to each vertex in the mesh (see Methods for full

details), we computed the association between the magnitude of cardiac motion and the 4D *survival* network's predicted risk score, yielding a set of regression coefficients (one per vertex) that were then mapped onto a template RV mesh, producing a 3D saliency map (Figure 3B). These show the contribution from spatially distant but functionally synergistic regions of the RV in influencing survival in PH.

Discussion

Machine learning algorithms have been used in a variety of motion analysis tasks from classifying complex traits to predicting future events from a given scene.^{9–11} We show that compressed representations of a dynamic biological system moving in 3D space offer a powerful approach for time-to-event analysis. In this example we demonstrate the effectiveness of a deep learning algorithm, trained to find correspondences between heart motion and patient outcomes, for efficiently predicting human survival.

The traditional paradigm of epidemiological research is to draw insight from large-scale clinical studies through linear regression modelling of conventional explanatory variables, but this approach does not embrace the dynamic physiological complexity of heart disease.¹² Even objective quantification of heart function by conventional analysis of cardiac imaging relies on crude measures of global contraction that are only moderately reproducible and insensitive to the underlying disturbances of cardiovascular physiology.¹³ Integrative approaches to risk classification have used unsupervised clustering of broad clinical variables to identify heart failure patients with distinct risk profiles,^{14, 15} while supervised machine learning algorithms can diagnose, risk stratify and predict adverse events from health record and registry data.^{16–18} In the wider health domain deep learning has achieved successes in forecasting survival from high-dimensional inputs such as cancer genomic profiles and gene expression data,^{19, 20} and in formulating personalised treatment recommendations.²¹

With the exception of natural image tasks, such as classification of skin lesions,²² biomedical imaging poses a number of challenges for machine learning as the datasets are often of limited scale, inconsistently annotated, and typically high-dimensional.²³ Architectures predominantly based on convolutional neural nets (CNNs), often using data augmentation strategies, have been successfully applied in computer vision tasks to enhance clinical images, segment organs and classify lesions.^{24, 25} Segmentation of cardiac images in the time domain is a well-established visual correspondence task that has recently achieved expert-level performance with FCN architectures.²⁶ Atlas-based analyses of cardiac geometry have demonstrated their value in disease classification and visualisation.^{27–29} Supervised principal components analysis of semi-automated segmentations has shown prognostic utility compared to conventional parameters,³⁰ but requires human selection of anatomical features and relies on simple pre-defined motion characteristics. In this work we harness the power of deep learning for both automated image analysis and inference - learning features predictive of survival from 3D cardiac motion using non-linear data transformations.

Autoencoding is a dimensionality reduction technique in which an encoder takes an input and maps it to a latent representation (lower-dimensional space) which is in turn mapped back to the space of the original input. The latter step represents an attempt to ‘reconstruct’ the input from the compressed (latent) representation, and this is done in such a way as to minimise the reconstruction error, i.e. the degree of discrepancy between the input and its reconstructed version. Our algorithm is based on a denoising autoencoder (DAE), a type of autoencoder which aims to extract more robust latent representations by corrupting the input with stochastic noise.³¹ While conventional autoencoders are used for unsupervised learning tasks we extend recent proposals for supervised autoencoders in which the learned representations are both reconstructive *and* discriminative.^{32–38} We achieved this by adding a prediction branch to the network with a loss function for survival inspired by the Cox proportional hazards model. A hybrid loss function, optimising the trade-off between survival prediction and accurate input reconstruction, is calibrated during training. The compressed representations of 3D motion predict survival more accurately than a composite measure of conventional manually-derived parameters measured on the same images and the improvement in performance is independent of clinical risk factors.

The main limitation of our study is relying on internal validation to evaluate predictive performance, and so the next step towards implementation is to train on larger and more diverse multicentre patient groups using image data and other prognostic variables, before performing external validation of survival prediction in a clinical setting against a benchmark of established risk prediction scores.³⁹ Autoencoders may be more prone to over-fitting than methods such as principal components analysis and are more computationally expensive to train. We mitigated over-fitting using dropout and L_1 regularization, and reduced the input space by down-sampling spatially-correlated data. We used routinely-acquired clinical data and applied normalisation to compare motion acquired at different temporal resolutions. Improvement in performance may be achievable at higher temporal resolutions, but would also increase the dimension of the input data. CMR provides accurate assessment of cardiac function but other imaging modalities may offer complementary prognostic markers.⁴⁰ Further enhancement in predictive performance may be achievable by modelling multiple observations over time, for instance using long short-term memory (LSTM) and other recurrent neural network architectures,^{41, 42} and handling independent competing risks.⁴³

Our approach enables fully automated and interpretable predictions of survival from moving clinical images - a task that has not been previously achieved in heart failure or other disease domains. This fast and scalable method is readily deployable and could have a substantial impact on clinical decision making and understanding of disease mechanisms. Extending this approach to other conditions where motion is predictive of survival is only constrained by the availability of suitable training cases with known outcomes.

Methods

Study Population

In a single-centre observational study, we analysed data collected from patients referred to the National Pulmonary Hypertension Service at the Imperial College Healthcare NHS Trust

between May 2004 and October 2017. The study was approved by the Heath Research Authority and all participants gave written informed consent. Criteria for inclusion were a documented diagnosis of Group 4 PH investigated by right heart catheterization (RHC) with a mean pulmonary artery pressure (mPAP) ≥ 25 mmHg and pulmonary capillary wedge pressure (PCWP) <15 mmHg; and signs of chronic thrombo-embolic disease present on either ventilation-perfusion scintigraphy or computed tomography pulmonary angiography. 44 All patients were treated in accordance with current guidelines including medical and surgical therapy as clinically indicated.⁵

MR Image Acquisition, Processing and Computational Image Analysis

The CMR protocol has been previously described in detail.³⁰ Briefly, imaging was performed on a 1.5 T Achieva (Philips, Best, Netherlands), using a standard clinical protocol based on international guidelines.⁴⁵ The specific images analysed in this study were retrospectively-gated cine sequences, in the short axis plane of the heart, with a reconstructed spatial resolution of $1.3 \times 1.3 \times 10.0$ mm and a typical temporal resolution of 29 ms. Images were stored on an open source data management system.⁴⁶ Manual volumetric analysis of the images was independently performed by accredited physicians using proprietary software (cmr42, Circle Cardiovascular Imaging, Calgary, Canada) according to international guidelines with access to all available images for each subject and no analysis time constraint.⁴⁷ The derived parameters included the strongest and most well-established volumetric and functional CMR findings for prognostication reported in disease-specific meta-analyses.^{48, 49}

We developed a CNN combined with image registration for shape-based biventricular segmentation of the CMR images. The pipeline method has three main components: segmentation, landmark localisation and shape registration. Firstly, a 2.5D multi-task FCN is trained to effectively and simultaneously learn segmentation maps and landmark locations from manually labelled volumetric CMR images. Secondly, multiple high-resolution 3D atlas shapes are propagated onto the network segmentation to form a smooth segmentation model. This step effectively induces a hard anatomical shape constraint and is fully automatic due to the use of predicted landmarks from the network.

We treat the problem of predicting segmentations and landmark locations as a multi-task classification problem. First, let us formulate the learning problem as follows: we denote the input training dataset by $S = \{(U_i, R_i, L_i), i = 1, \dots, N_t\}$, where N_t is the sample size of the training data, $U_i = \{u_j^i, j = 1, \dots, |U_i|\}$ is the raw input CMR volume, $R_i = \{r_j^i, j = 1, \dots, |R_i|\}$, $r_j^i \in \{1, \dots, N_r\}$ are the ground truth region labels for volume U_i ($N_r = 5$ representing 4 regions and background), and $L_i = \{l_j^i, j = 1, \dots, |L_i|\}$, $l_j^i \in \{1, \dots, N_l\}$ are the labels representing ground truth landmark locations for U_i ($N_l = 7$ representing 6 landmark locations and background). Note that $|U_i| = |R_i| = |L_i|$ stands for the total number of voxels in a CMR volume. Let \mathbf{W} denote the set of all network layer parameters. In a supervised setting, we minimise the following objective function via standard (backpropagation) stochastic gradient descent (SGD):

$$L(\mathbf{W}) = L_S(\mathbf{W}) + \alpha L_D(\mathbf{W}) + \beta L_L(\mathbf{W}) + \gamma \|\mathbf{W}\|_F^2, \quad (1)$$

where α , β and γ are weight coefficients balancing the four terms. $L_S(\mathbf{W})$ and $L_D(\mathbf{W})$ are the region-associated losses that enable the network to predict segmentation maps. $L_L(\mathbf{W})$ is the landmark-associated loss for predicting landmark locations. $\|\mathbf{W}\|_F^2$, known as the weight decay term, represents the Frobenius norm on the weights \mathbf{W} . This term is used to prevent the network from overfitting. The training problem is therefore to estimate the parameters \mathbf{W} associated with all the convolutional layers. By minimising (1), the network is able to simultaneously predict segmentation maps and landmark locations. The definitions of the loss functions $L_S(\mathbf{W})$, $L_D(\mathbf{W})$ and $L_L(\mathbf{W})$, used for predicting landmarks and segmentation labels, have been described previously.⁵⁰

The FCN segmentations are used to perform a non-rigid registration using cardiac atlases built from >1000 high resolution images,⁵¹ allowing shape constraints to be inferred. This approach produces accurate, high-resolution and anatomically smooth segmentation results from input images with low through-slice resolution thus preserving clinically-important global anatomical features. The data were split in the ratio 70:30 for training and evaluation respectively. Motion tracking was performed for each subject using a 4D spatio-temporal B-spline image registration method with a sparseness regularisation term.⁵² The motion field estimate is represented by a displacement vector at each voxel and at each time frame $t = 1, \dots, 20$. Temporal normalisation was performed before motion estimation to ensure consistency across the cardiac cycle.

Spatial normalisation of each patient's data was achieved by registering the motion fields to a template space. A template image was built by registering the high-resolution atlases at the end-diastolic frame and then computing an average intensity image. In addition, the corresponding ground-truth segmentations for these high-resolution images were averaged to form a segmentation of the template image. A template surface mesh was then reconstructed from its segmentation using a 3D surface reconstruction algorithm. The motion field estimate lies within the reference space of each subject and so to enable inter-subject comparison all the segmentations were aligned to this template space by non-rigid B-spline image registration.⁵³ We then warped the template mesh using the resulting non-rigid deformation and mapped it back to the template space. Twenty surface meshes, one for each temporal frame, were subsequently generated by applying the estimated motion fields to the warped template mesh accordingly. Consequently, the surface mesh of each subject at each frame contained the same number of vertices (18, 028) which maintained their anatomical correspondence across temporal frames, and across subjects (Figure 5).

Characterization of right ventricular motion

The time-resolved 3D meshes described in the previous section were used to produce a relevant representation of cardiac motion - in this example of right-side heart failure limited to the RV. For this purpose, we utilized a sparser version of the meshes (down-sampled by a factor of ~90) with 202 vertices. Anatomical correspondence was preserved in this process

by utilizing the same vertices across all meshes. To characterize motion, we adapted an approach outlined in Bai *et al* (2015).⁵⁴

This approach is used to produce a simple numerical representation of the trajectory of each vertex, i.e. the path each vertex traces through space during a cardiac cycle (see Figure 1B). Let (x_{vt}, y_{vt}, z_{vt}) represent the Cartesian coordinates of vertex v ($v = 1, \dots, 202$) at the t^{th} time frame ($t = 1, \dots, 20$) of the cardiac cycle. At each time frame $t = 2, 3, \dots, 20$, we compute the coordinate-wise displacement of each vertex from its position at time frame 1. This yields the following one-dimensional input vector:

$$\mathbf{x} = (x_{vt} - x_{v1}, y_{vt} - y_{v1}, z_{vt} - z_{v1})_{\substack{1 \leq v \leq 202 \\ 2 \leq t \leq 20}} \quad (2)$$

Vector \mathbf{x} has length 11,514 ($3 \times 19 \times 202$), and was used as the input feature for our prediction network.

Network Design and Training

Our 4D*survival* network structure is summarized in Figure 6. We aimed to produce an architecture capable of learning a low-dimensional representation of RV motion that robustly captures prognostic features indicative of poor survival. The architecture's hybrid design combines a denoising autoencoder,⁵⁵ with a Cox proportional hazards model (described below).⁵⁶

As before, we denote our input vector by $\mathbf{x} \in \mathbb{R}^{d_p}$, where $d_p = 11,514$, the input dimensionality. Our network is based on a DAE, an autoencoder variant which learns features robust to noise.⁵⁵ The input vector \mathbf{x} feeds directly into the encoder, the first layer of which is a stochastic masking filter that produces a corrupted version of \mathbf{x} . The masking is implemented using random dropout,⁵⁷ i.e. we randomly set a fraction m of the elements of vector \mathbf{x} to zero (the value of m is treated as an optimizable network parameter). The corrupted input from the masking filter is then fed into a hidden layer, the output of which is in turn fed into a central layer. This central layer represents the latent code, i.e. the encoded/compressed representation of the input. This central layer is referred to as the 'code', or 'bottleneck' layer. Therefore we may consider the encoder as a function $\phi(\cdot)$ mapping the input $\mathbf{x} \in \mathbb{R}^{d_p}$ to a latent code $\phi(\mathbf{x}) \in \mathbb{R}^{d_h}$, where $d_h \ll d_p$ (for notational convenience we consider the corruption step as part of the encoder). This produces a compressed representation whose dimensionality is much lower than that of the input (an undercomplete representation).⁵⁸ Note that the number of units in the encoder's hidden layer, and the dimensionality of the latent code (d_h) are not predetermined but, rather, treated as optimisable network parameters. The latent code $\phi(\mathbf{x})$ is then fed into the second component of the DAE, a multilayer decoder network that upsamples the code back to the original input dimension d_p . Like the encoder, the decoder has one intermediate hidden layer that feeds into the final layer, which in turn outputs a decoded representation (with dimension d_p matching that of the input). The size of the decoder's intermediate hidden layer is constrained to match that of the encoder network, to give the autoencoder a symmetric architecture. Dissimilarity between the original (uncorrupted) input \mathbf{x} and the decoder's

reconstructed version (denoted here by $\psi(\phi(x))$) is penalized by minimizing a loss function of general form $L(x, \psi(\phi(x)))$. Herein, we chose a simple mean squared error form for L :

$$L_r = \frac{1}{n} \sum_{i=1}^n \|x_i - \psi(\phi(x_i))\|^2, \quad (3)$$

where n represents the sample size. Minimizing this loss forces the autoencoder to reconstruct the input from a corrupted/incomplete version, thereby facilitating the generation of a latent representation with robust features. Further, to ensure that these learned features are actually relevant for survival prediction, we augmented the autoencoder network by adding a prediction branch. The latent representation learned by the encoder $\phi(x)$ is therefore linked to a linear predictor of survival (see equation 4 below), in addition to the decoder. This encourages the latent representation $\phi(x)$ to contain features which are simultaneously robust to noisy input and salient for survival prediction. The prediction branch of the network is trained with observed outcome data, i.e. survival/follow-up time. For each subject, this is time elapsed from MRI acquisition until death (all-cause mortality), or if the subject is still alive, the last date of follow-up. Also, patients receiving surgical interventions were censored at the date of surgery. This type of outcome is called a *right-censored time-to-event* outcome,⁵⁹ and is typically handled using survival analysis techniques, the most popular of which is Cox's proportional hazards regression model:⁵⁶

$$\log \frac{h_i(t)}{h_0(t)} = \beta_1 z_{i1} + \beta_2 z_{i2} + \dots + \beta_p z_{ip} \quad (4)$$

Here, $h_i(t)$ represents the hazard function for subject i , i.e the 'chance' (normalized probability) of subject i dying at time t . The term $h_0(t)$ is a baseline hazard level to which all subject-specific hazards $h_i(t)$ ($i = 1, \dots, n$) are compared. The key assumption of the Cox survival model is that the *hazard ratio* $h_i(t)/h_0(t)$ is constant with respect to time (*proportional hazards* assumption).⁵⁶ The natural logarithm of this ratio is modeled as a weighted sum of a number of predictor variables (denoted here by z_{i1}, \dots, z_{ip}), where the weights/coefficients are unknown parameters denoted by β_1, \dots, β_p . These parameters are estimated via maximization of the Cox proportional hazards partial likelihood function:

$$\log \mathcal{L}(\boldsymbol{\beta}) = \sum_{i=1}^n \delta_i \left\{ \boldsymbol{\beta}' z_i - \log \sum_{j \in R(t_i)} e^{\boldsymbol{\beta}' z_j} \right\} \quad (5)$$

In the expression above, z_j is the vector of predictor/explanatory variables for subject i , δ_j is an indicator of subject i 's status (0=Alive, 1=Dead) and $R(t_j)$ represents subject i 's risk set, i.e. subjects still alive (and thus at risk) at the time subject i died or became censored ($\{j: t_j > t_i\}$).

We adapt this loss function for our neural network architecture as follows:

$$L_s = - \sum_{i=1}^n \delta_i \left\{ \mathbf{W}' \phi(\mathbf{x}_i) - \log \sum_{j \in R(t_i)} e^{\mathbf{W}' \phi(\mathbf{x}_j)} \right\} \quad (6)$$

The term \mathbf{W}' denotes a $(1 \times d_h)$ vector of weights, which when multiplied by the d_h -dimensional latent code $\phi(\mathbf{x})$ yields a single scalar ($\mathbf{W}' \phi(\mathbf{x}_i)$) representing the survival prediction (specifically, natural logarithm of the hazard ratio) for subject i . Note that this makes the prediction branch of our 4D*survival* network essentially a simple linear Cox proportional hazards model, and the predicted output may be seen as an estimate of the log hazard ratio (see Equation 4).

For our network, we combine this survival loss with the reconstruction loss from equation 3 to form a hybrid loss given by:

$$L_{\text{hybrid}} = \alpha L_r + \gamma L_s = \alpha \left[\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \psi(\phi(\mathbf{x}_i))\|^2 \right] + \gamma \left[- \sum_{i=1}^n \delta_i \left\{ \mathbf{W}' \phi(\mathbf{x}_i) - \log \sum_{j \in R(t_i)} e^{\mathbf{W}' \phi(\mathbf{x}_j)} \right\} \right] \quad (7)$$

The terms α and γ are used to calibrate the contributions of each term to the overall loss, i.e. to control the tradeoff between survival prediction versus accurate input reconstruction. During network training, they are treated as optimisable network hyperparameters, with γ chosen to equal $1 - \alpha$ for convenience.

The loss function was minimized via backpropagation. To avoid overfitting and to encourage sparsity in the encoded representation, we applied L_1 regularization. The rectified linear unit (ReLU) activation function was used for all layers, except the prediction output layer (linear activation was used for this layer). Using the adaptive moment estimation (*Adam*) algorithm, the network was trained for 100 epochs with a batch size of 16 subjects. The learning rate is treated as a hyperparameter (see Table 2). During training, the random dropout (input corruption) was repeated at every backpropagation pass. The network was implemented and trained in the Python deep learning libraries *TensorFlow*⁶⁰ and *Keras*,⁶¹ on a high-performance computing cluster with an Intel Xeon E5-1660 CPU and NVIDIA TITAN Xp GPU. The entire training process (including hyperparameter search and bootstrap-based internal validation [see subsections below]) took a total of 131 hours.

Hyperparameter Tuning

To determine optimal hyperparameter values, we utilized particle swarm optimization (PSO),⁶² a gradient-free meta-heuristic approach to finding optima of a given objective function. Inspired by the social foraging behavior of birds, PSO is based on the principle of *swarm intelligence*, which refers to problem-solving ability that arises from the interactions of simple information-processing units.⁶³ In the context of hyperparameter tuning, it can be

used to maximize the prediction accuracy of a model with respect to a set of potential hyperparameters.⁶⁴ We used PSO to choose the optimal set of hyperparameters from among predefined ranges of values (summarized in Table 2). We ran the PSO algorithm for 50 iterations, at each step evaluating candidate hyperparameter configurations using 6-fold cross-validation. The hyperparameters at the final iteration were chosen as the optimal set. This procedure was implemented via the Python library *Optunity*.⁶⁵

Model Validation and Comparison

Predictive Accuracy Metric—Discrimination was evaluated using Harrell’s concordance index,⁶⁶ an extension of area under the receiver operating characteristic curve (AUC) to censored time-to-event data:

$$C = \frac{\sum_{i,j} \delta_i \cdot I(\eta_i > \eta_j) \cdot I(t_i < t_j)}{\sum_{i,j} \delta_i \cdot I(t_i < t_j)} \quad (8)$$

In the above equation, the indices i and j refer to pairs of subjects in the sample and $I()$ denotes an indicator function that evaluates to 1 if its argument is true (and 0 otherwise). Symbols η_i and η_j denote the predicted risks for subjects i and j . The numerator tallies the number of subject pairs (i, j) where the pair member with greater predicted risk has shorter survival, representing agreement (concordance) between the model’s risk predictions and ground-truth survival outcomes. Multiplication by δ_i restricts the sum to subject pairs where it is possible to determine who died first (i.e. informative pairs). The C index therefore represents the fraction of informative pairs exhibiting concordance between predictions and outcomes. In this sense, the index has a similar interpretation to the AUC (and consequently, the same range).

Internal Validation—In order to get a sense of how well our model would generalize to an external validation cohort, we assessed its predictive accuracy within the training sample using a bootstrap-based procedure recommended in the guidelines for *Transparent Reporting of a multivariable model for Individual Prognosis Or Diagnosis (TRIPOD)*.⁶⁷ This procedure attempts to derive realistic, ‘optimism-adjusted’ estimates of the model’s generalization accuracy using the training sample.⁶⁸ Below, we outline the steps of the procedure:

- (Step 1)** A prediction model was developed on the full training sample (size n), utilizing the hyperparameter search procedure discussed above to determine the best set of hyperparameters. Using the optimal hyperparameters, a final model was trained on the full sample. Then the Harrell’s concordance index (C) of this model was computed on the full sample, yielding the *apparent* accuracy, i.e. the inflated accuracy obtained when a model is tested on the same sample on which it was trained/optimized.
- (Step 2)** A bootstrap sample was generated by carrying out n random selections (with replacement) from the full sample. On this bootstrap sample, we developed a model (applying exactly the same training and hyperparameter

search procedure used in Step 1) and computed C for the bootstrap sample (henceforth referred to as *bootstrap performance*). Then the performance of this bootstrap-derived model on the original data (the full training sample) was also computed (henceforth referred to as *test performance*)

- (Step 3) For each bootstrap sample, the optimism was computed as the difference between the bootstrap performance and the test performance.
- (Step 4) Steps 2-3 were repeated B times (where $B=100$).
- (Step 5) The optimism estimates derived from Steps 2-4 were averaged across the B bootstrap samples and the resulting quantity was subtracted from the apparent predictive accuracy from Step 1.

This procedure yields an optimism-corrected estimate of the model's concordance index:

$$C_{corrected} = C_{full}^{full} - \frac{1}{B} \sum_{b=1}^B (C_b^b - C_b^{full}) \quad (9)$$

Above, symbol $C_{s_1}^{s_2}$ refers to the concordance index of a model trained on sample s_1 and tested on sample s_2 . The first term refers to the *apparent* predictive accuracy, i.e. the (inflated) concordance index obtained when a model trained on the full sample is then tested on the same sample. The second term is the average *optimism* (difference between *bootstrap performance* and *test performance*) over the B bootstrap samples. It has been demonstrated that this sample-based average is a nearly unbiased estimate of the expected value of the optimism that would be observed in external validation.^{68–71} Subtraction of this optimism estimate from the apparent predictive accuracy gives the optimism-corrected predictive accuracy.

Conventional Parameter model—As a benchmark comparison to our RV motion model, we trained a Cox proportional hazards model using conventional RV volumetric indices including RV end-diastolic volume (RVEDV), RV end-systolic volume (RVESV) and the difference between these measures expressed as a percentage of RVEDV, RV ejection fraction (RVEF) as survival predictors. We also trained a model on strain-related measures of mechanical function with tensors in the longitudinal, radial and circumferential directions.⁷² A last model was trained on both the CMR parameters and a set of clinical risk factors,⁷³ that comprised age, sex, six minute walk distance, functional class and mPAP using the *missForest* algorithm to impute any missing values.⁷⁴ To account for collinearity among these predictor variables, a regularization term was added to the Cox partial likelihood function:

$$\log L(\beta) = \sum_{i=1}^n \delta_i \left\{ \beta' x_i - \log \sum_{j \in R(t_i)} e^{\beta' x_j} \right\} + \frac{1}{2} \lambda \| \beta \|^2 \quad (10)$$

In the equation above, λ is a parameter that controls the strength of the penalty. The optimal value of λ was selected via cross-validation. Internal validation of these models was carried out using the bootstrap-based procedure outlined in the previous section. Model comparisons were carried out using the R package *survcomp75* to compare concordance index measures (see supplementary materials for further details).

Model Interpretation—To facilitate interpretation of our 4D*survival* network we used Laplacian Eigenmaps to project the learned latent code into two dimensions,⁸ allowing latent space visualization. Neural networks derive predictions through multiple layers of nonlinear transformations on the input data. This complex architecture does not lend itself to straightforward assessment of the relative importance of individual input features. To tackle this problem we used a simple regression-based inferential mechanism to evaluate the contribution of motion in various regions of the RV to the model’s predicted risk. For each of the 202 vertices in our RV mesh models we computed a single summary measure of motion by averaging the displacement magnitudes across 19 frames. This yielded one mean displacement value per vertex. This process was repeated across all subjects. Then we regressed the predicted risk scores onto these vertex-wise mean displacement magnitude measures using a mass univariate approach, i.e. for each vertex v ($v = 1, \dots, 202$), we fitted a linear regression model where the dependent variable was predicted risk score, and the independent variable was average displacement magnitude of vertex v . Each of these 202 univariate regression models was fitted on all subjects and yielded one regression coefficient representing the effect of motion at a vertex on predicted risk. The absolute values of these coefficients, across all vertices, were then mapped onto a template RV mesh to provide a visualization of the differential contribution of various anatomical regions to predicted risk.

Data and code availability

Algorithms, motion models and statistical analysis are publicly available on *GitHub* under a GNU General Public License.⁷⁶ A training simulation is available as a *Docker* image with an interactive *Jupyter* notebook hosted on *Code Ocean*.⁷⁷ Personal data are not available due to privacy restrictions.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The research was supported by the British Heart Foundation (NH/17/1/32725, RE/13/4/30184); National Institute for Health Research (NIHR) Biomedical Research Centre based at Imperial College Healthcare NHS Trust and Imperial College London; and the Medical Research Council, UK. The TITAN Xp GPU used for this research was kindly donated by the NVIDIA Corporation.

References

1. Wang, L, Zhao, G, Cheng, L, Pietikäinen, M. Machine learning for vision-based motion analysis: Theory and techniques. Springer; 2010.
2. Mei, T; Zhang, C. Deep learning for intelligent video analysis. 2017. URL <https://www.microsoft.com/en-us/research/publication/deep-learning-intelligent-video-analysis/>

3. Liang F, Xie W, Yu Y. Beating heart motion accurate prediction method based on interactive multiple model: An information fusion approach. *BioMed Res Int.* 2017; 2017:9.
4. Savarese G, Lund LH. Global public health burden of heart failure. *Cardiac Fail Rev.* 2017; 3:7–11.
5. Galie N, et al. 2015 ESC/ERS guidelines for the diagnosis and treatment of pulmonary hypertension: The Joint Task Force for the Diagnosis and Treatment of Pulmonary Hypertension of the European Society of Cardiology (ESC) and the European Respiratory Society (ERS): Endorsed by: Association for European Paediatric and Congenital Cardiology (AEPC), International Society for Heart and Lung Transplantation (ISHLT). *Eur Hear J.* 2016; 37:67–119.
6. Puyol-Antón E, et al. A multimodal spatiotemporal cardiac motion atlas from MR and ultrasound data. *Med Image Anal.* 2017; 40:96–110. [PubMed: 28646674]
7. Scatteia A, Baritussio A, Bucciarelli-Ducci C. Strain imaging using cardiac magnetic resonance. *Hear Fail Rev.* 2017; 22:465–476.
8. Belkin, M, Niyogi, P. Laplacian eigenmaps and spectral techniques for embedding and clustering *Advances in Neural Information Processing Systems*. Dietterich, TG, Becker, S, Ghahramani, Z, editors Vol. 14. MIT Press; 2002. 585–591.
9. Li K, Javer A, Keaveny EE, Brown AEX. Recurrent neural networks with interpretable cells predict and classify worm behaviour. *bioRxiv.* 2017
10. Walker J, Doersch C, Gupta A, Hebert M. An Uncertain Future: Forecasting from Static Images using Variational Autoencoders. *ArXiv 1606.07873.* 2016
11. Bütepage J, Black M, Kragic D, Kjellström H. Deep representation learning for human motion prediction and classification. *ArXiv 1702.07486.* 2017
12. Johnson KW, et al. Enabling precision cardiology through multiscale biology and systems medicine. *JACC: Basic to Transl Sci.* 2017; 2:311–327.
13. Cikes M, Solomon SD. Beyond ejection fraction: an integrative approach for assessment of cardiac structure and function in heart failure. *Eur Hear J.* 2016; 37:1642–50.
14. Ahmad T, et al. Clinical implications of chronic heart failure phenotypes defined by cluster analysis. *J Am Coll Cardiol.* 2014; 64:1765–74. [PubMed: 25443696]
15. Shah SJ, et al. Phenomapping for novel classification of heart failure with preserved ejection fraction. *Circ.* 2015; 131:269–79.
16. Awan SE, Sohail F, Sanfilippo FM, Bennamoun M, Dwivedi G. Machine learning in heart failure: ready for prime time. *Curr Opin Cardiol.* 2018; 33:190–195. [PubMed: 29194052]
17. Tripoliti EE, Papadopoulos TG, Karanasiou GS, Naka KK, Fotiadis DI. Heart failure: Diagnosis, severity estimation and prediction of adverse events through machine learning techniques. *Comput Struct Biotechnol J.* 2017; 15:26–47. [PubMed: 27942354]
18. Ambale-Venkatesh B, et al. Cardiovascular event prediction by machine learning: The multi-ethnic study of atherosclerosis. *Circ Res.* 2017; 121:1092–1101. [PubMed: 28794054]
19. Yousefi S, et al. Predicting clinical outcomes from large scale cancer genomic profiles with deep survival models. *Sci Rep.* 2017; 7
20. Ching T, Zhu X, Garmire LX. Cox-nnet: An artificial neural network method for prognosis prediction of high-throughput omics data. *PLOS Comput Biol.* 2018; 14:1–18.
21. Katzman J, et al. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Med Res Meth.* 2018; 18:1–12.
22. Esteva A, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nat.* 2017; 542:115–118.
23. Ching T, et al. Opportunities and obstacles for deep learning in biology and medicine. *J R Soc Interface.* 2018; 15
24. Litjens G, et al. A survey on deep learning in medical image analysis. *Med Image Anal.* 2017; 42:60–88. [PubMed: 28778026]
25. Shen D, Wu G, Suk HI. Deep learning in medical image analysis. *Annu Rev Biomed Eng.* 2017; 19:221–248. [PubMed: 28301734]
26. Bai W, et al. Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *J Cardiovasc Magn Reson.* 2018; 20:65. [PubMed: 30217194]

27. Piras P, et al. Morphologically normalized left ventricular motion indicators from MRI feature tracking characterize myocardial infarction. *Sci Rep.* 2017; 7
28. Zhang X, et al. Orthogonal decomposition of left ventricular remodeling in myocardial infarction. *Gigascience.* 2017; 6:1–15.
29. Zhang X, et al. Atlas-based quantification of cardiac remodeling due to myocardial infarction. *PLoS One.* 2014; 9
30. Dawes T, et al. Machine learning of three-dimensional right ventricular motion enables outcome prediction in pulmonary hypertension: a cardiac MR imaging study. *Radiol.* 2017; 283:381–390.
31. Rifai, S; Vincent, P; Muller, X; Glorot, X; Bengio, Y. Contractive auto-encoders: Explicit invariance during feature extraction. *Proceedings of the 28th International Conference on International Conference on Machine Learning; Omnipress; 2011.* 833–840.
32. Rolfe JT, LeCun Y. Discriminative Recurrent Sparse Auto-Encoders. *ArXiv 1301.3775.* 2013
33. Huang R, Liu C, Li G, Zhou J. Adaptive deep supervised autoencoder based image reconstruction for face recognition. *Math Probl Eng.* 2016; 2016:14.
34. Du F, Zhang J, Ji N, Hu J, Zhang C. Discriminative representation learning with supervised auto-encoder. *Neur Proc Lett.* 2018
35. Zaghbani S, Boujneh N, Bouhlel MS. Age estimation using deep learning. *Comp Elec Eng.* 2018; 68:337–347.
36. Beaulieu-Jones BK, Greene CS. Semi-supervised learning of the electronic health record for phenotype stratification. *J Biomed Inf.* 2016; 64:168–178.
37. Shakeri, M; Lombaert, H; Tripathi, S; Kadoury, S. Deep spectral-based shape features for alzheimer's disease classification. In: Reuter, M; Wachinger, C; Lombaert, H, editors *International Workshop on Spectral and Shape Analysis in Medical Imaging; Cham: Springer; 2016.* 15–24.
38. Biffi, C; , et al. Learning interpretable anatomical features through deep generative models: Application to cardiac remodeling. *International Conference on Medical Image Computing and Computer-Assisted Intervention; Springer; 2018.*
39. Dawes TJW, Bello GA, O'Regan DP. Multicentre study of machine learning to predict survival in pulmonary hypertension. *Open Sci Framew.* 2018; doi: 10.17605/OSF.IO/BG6T9
40. Grapsa J, et al. Echocardiographic and hemodynamic predictors of survival in precapillary pulmonary hypertension: Seven-year follow-up. *Circ Cardiovasc Imaging.* 2015; 8
41. Bao W, Yue J, Rao Y. A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLoS ONE.* 2017; 12
42. Lim B, van der Schaar M. Disease-atlas: Navigating disease trajectories with deep learning. *ArXiv 1803.10254.* 2018
43. Lee C, Zame WR, Yoon J, van der Schaar M. DeepHit: A deep learning approach to survival analysis with competing risks. *AAAI.* 2018
44. Gopalan D, Delcroix M, Held M. Diagnosis of chronic thromboembolic pulmonary hypertension. *Eur Respir Rev.* 2017; 26
45. Kramer C, Barkhausen J, Flamm S, Kim R, Nagel E. Society for cardiovascular magnetic resonance board of trustees task force on standardized protocols. standardized cardiovascular magnetic resonance (CMR) protocols 2013 update. *J Cardiovasc Magn Reson.* 2013; 15
46. Woodbridge M, Fagiolo G, O'Regan DP. MRIdb: medical image management for biobank research. *J Digit Imaging.* 2013; 26:886–890. [PubMed: 23619930]
47. Schulz-Menger J, et al. Standardized image interpretation and post processing in cardiovascular magnetic resonance: Society for cardiovascular magnetic resonance (SCMR) board of trustees task force on standardized post processing. *J Cardiovasc Magn Reson.* 2013; 15:35. [PubMed: 23634753]
48. Baggen VJ, et al. Cardiac magnetic resonance findings predicting mortality in patients with pulmonary arterial hypertension: a systematic review and meta-analysis. *Eur Radiol.* 2016; 26:3771–3780. [PubMed: 26847041]
49. Hulshof HG, et al. Prognostic value of right ventricular longitudinal strain in patients with pulmonary hypertension: a systematic review and meta-analysis. *Eur Hear J Cardiovasc Imaging.* 2018

50. Duan J, et al. Automatic 3D bi-ventricular segmentation of cardiac images by a shape-constrained multi-task deep learning approach. ArXiv 1808.08578. 2018
51. Bai W, et al. A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion. *Med Imag Anal.* 2015; 26:133–145.
52. Shi W, et al. Temporal sparse free-form deformations. *Med Imag Anal.* 2013; 17:779–789.
53. Rueckert D, et al. Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Trans Med Imaging.* 1999; 18:712–721. [PubMed: 10534053]
54. Bai, W, , et al. Learning a Global Descriptor of Cardiac Motion from a Large Cohort of 1000+ Normal Subjects. Vol. 9126. Springer; Cham: 2015.
55. Vincent P, Larochelle H, Lajoie I, Bengio Y, Manzagol P-A. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J Mach Learn Res.* 2010; 11:3371–3408.
56. Cox D. Regression models and life-tables. *J Roy Stat Soc B.* 1972; 34:187–220.
57. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A simple way to prevent neural networks from overfitting. *J Mach Learn Res.* 2014; 15:1929–1958.
58. Goodfellow, I, Bengio, Y, Courville, A. *Deep Learning.* MIT Press; 2016. URL <http://www.deeplearningbook.org>
59. Faraggi D, Simon R. A neural network model for survival data. *Stat Med.* 1995; 14:73–82. [PubMed: 7701159]
60. Abadi, M; , et al. TensorFlow: Large-scale machine learning on heterogeneous systems. 2015. URL <https://www.tensorflow.org/>
61. Chollet, F; , et al. Keras. 2015. URL <https://keras.io>
62. Kennedy, J; Eberhart, R. Particle swarm optimization. *Proc IEEE Int Conf Neural Net;* 1995. 1942–1948.
63. Engelbrecht, A. *Fundamentals of computational swarm intelligence.* Wiley; Chichester: 2005.
64. Lorenzo, PR; Nalepa, J; Kawulok, M; Ramos, LS; Pastor, JR. Particle swarm optimization for hyper-parameter selection in deep neural networks. *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17;* 2017. 481–488.
65. Claesen M, Simm J, Popovic D, De Moor B. Hyperparameter tuning in Python using Optunity. *Proc Int Work on Tech Comput for Mach Learn Math Eng.* 2014; 9
66. Harrell F, Califf R, Pryor D, Lee K, Rosati R. Evaluating the yield of medical tests. *JAMA.* 1982; 247
67. Moons K, et al. Transparent reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): Explanation and elaboration. *Ann Intern Med.* 2015; 162:W1–W73. [PubMed: 25560730]
68. Harrell F, Lee K, Mark D. Tutorial in biostatistics: multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Stat Med.* 1996; 15:361–387. [PubMed: 8668867]
69. Efron B. Estimating the error rate of a prediction rule: some improvements on cross-validation. *JASA.* 1983; 78:316–331.
70. Efron, B, Tibshirani, R. Cross-validation and other estimates of prediction error An Introduction to the Bootstrap. Vol. chap. 17. Chapman & Hall; New York: 1993. 237–257.
71. Smith G, Seaman S, Wood A, Royston P, White I. Correcting for optimistic prediction in small data sets. *Am J Epidem.* 2014; 180
72. Liu B, et al. Normal values for myocardial deformation within the right heart measured by feature-tracking cardiovascular magnetic resonance imaging. *Int J Cardiol.* 2018; 252:220–223. DOI: 10.1016/j.ijcard.2017.10.106 [PubMed: 29133275]
73. Gall H, et al. The Giessen pulmonary hypertension registry: Survival in pulmonary hypertension subgroups. *J Hear Lung Transpl.* 2017; 36:957–967.
74. Stekhoven DJ, Buhlmann P. missForest—non-parametric missing value imputation for mixed-type data. *Bioinforma.* 2011; 28:112–118.

75. Schroder MS, Culhane AC, Quackenbush J, Haibe-Kains B. survcomp: an R/Bioconductor package for performance assessment and comparison of survival models. *Bioinforma.* 2011; 27:3206–3208.
76. Bello GA, et al. Deep learning cardiac motion analysis for human survival prediction (4Dsurvival). 2019; doi: 10.5281/zenodo.1451540
77. Bello GA, et al. Deep learning cardiac motion analysis for human survival prediction (4Dsurvival). 2019; doi: 10.24433/CO.8519672.v1

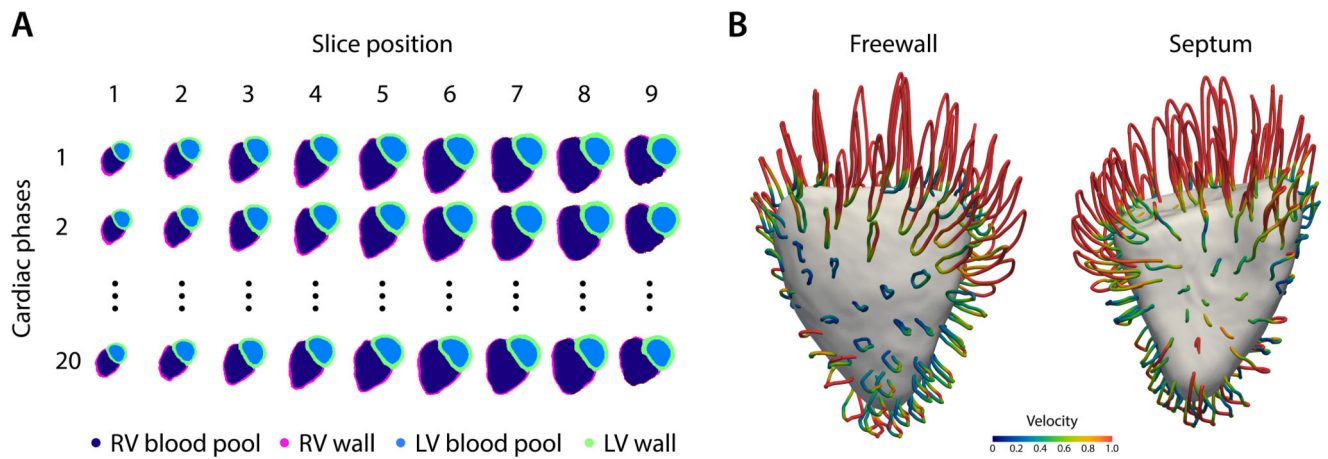


Figure 1.

A) An example of an automatic cardiac image segmentation of each short-axis cine image from apex (slice 1) to base (slice 9) across 20 temporal phases. Data were aligned to a common reference space to build a population model of cardiac motion. B) Trajectory of right ventricular contraction and relaxation averaged across the study population plotted as looped pathlines for a sub-sample of 100 points on the heart (magnification factor of x4). Colour represents relative myocardial velocity at each phase of the cardiac cycle. A surface-shaded model of the heart is shown at end-systole. These dense myocardial motion fields for each patient were used as an input to the prediction network. LV, left ventricular; RV, right ventricular.

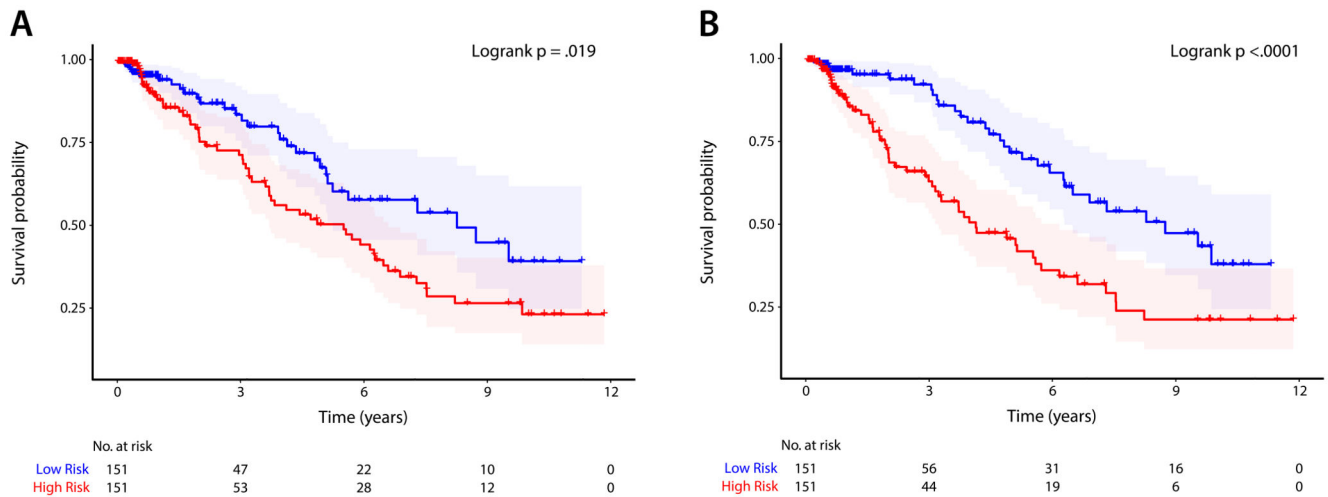


Figure 2.

Kaplan-Meier plots for A) a conventional parameter model using a composite of manually-derived volumetric measures, and B) a deep learning prediction model (4Dsurvival) whose input was time-resolved three dimensional models of cardiac motion. For both models, patients were divided into low- and high-risk groups by median risk score. Survival function estimates for each group (with 95% confidence intervals) are shown. For each plot, the Logrank test was performed to compare survival curves between risk groups (conventional parameter model: $\chi^2 = 5.5$, $p = .019$; 4Dsurvival: $\chi^2 = 15.6$, $p < .0001$)

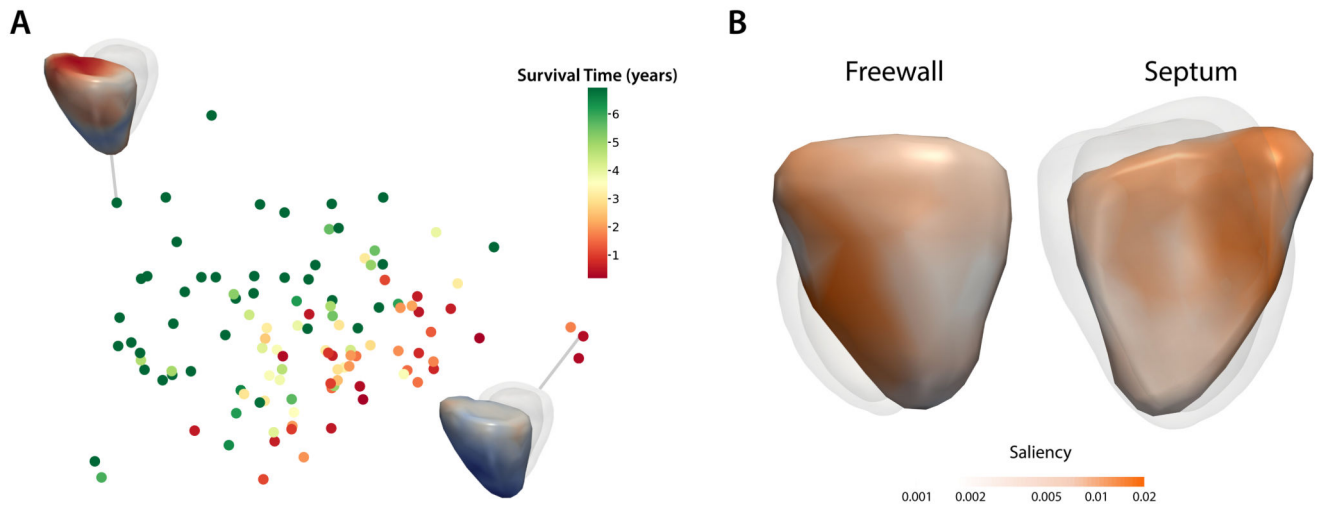


Figure 3.

A) A 2-dimensional projection of latent representations of cardiac motion in the 4Dsurvival network labelled by survival time. A visualisation of RV motion is shown for two patients with contrasting risks. B) Saliency map showing regional contributions to survival prediction by right ventricular motion. Absolute regression coefficients are expressed on a log-scale.

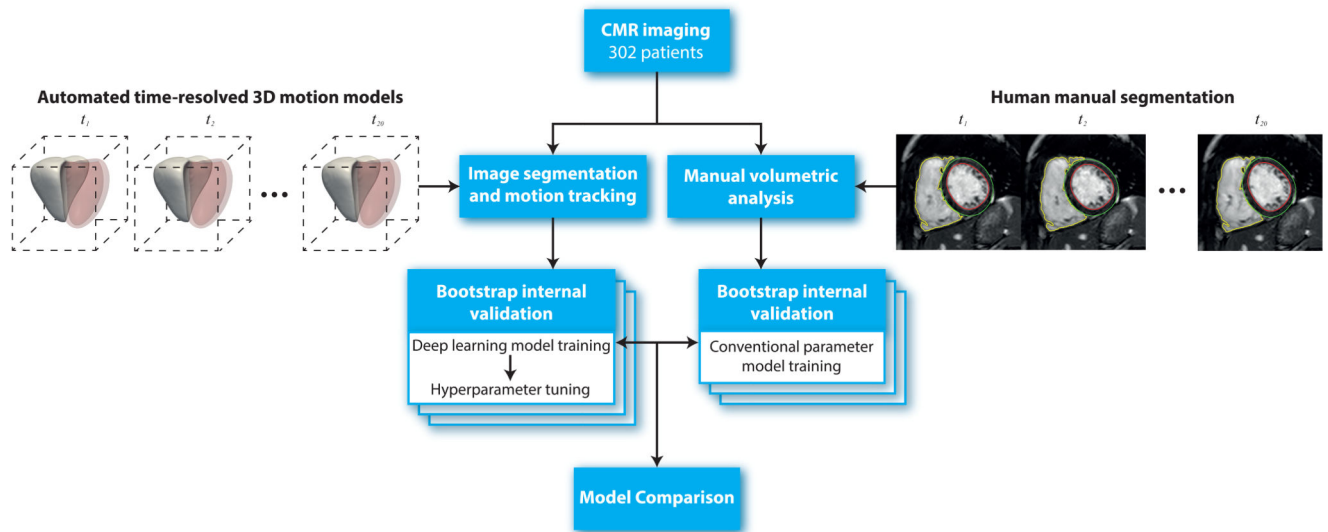


Figure 4.

Flowchart to show the design of the study. In total 302 patients with CMR imaging had both manual volumetric analysis and automated image segmentation (right ventricle shown in solid white, left ventricle in red) across 20 temporal phases ($t = 1, \dots, 20$). Internal validity of the predictive performance of a conventional parameter model and a deep learning motion model was assessed using bootstrapping. CMR, cardiac magnetic resonance.




Figure 5.

The architecture of the segmentation algorithm. A fully convolutional network takes each stack of cine images as an input, applies a branch of convolutions, learns image features from fine to coarse levels, concatenates multi-scale features and finally predicts the segmentation and landmark location probability maps simultaneously. These maps, together with the ground truth landmark locations and label maps, are then used in the loss function (see Equation 1) which is minimised via stochastic gradient descent.




Figure 6.

The prediction network (4Dsurvival) is a denoising autoencoder that takes time-resolved cardiac motion meshes as its input (right ventricle shown in solid white, left ventricle in red). For the sake of simplicity two hidden layers, one immediately preceding and the other immediately following the central layer (latent code layer), have been excluded from the diagram. The autoencoder learns a task-specific latent code representation trained on observed outcome data, yielding a latent representation optimised for survival prediction that is robust to noise. The actual number of latent factors is treated as an optimisable parameter.

Table 1

Patient characteristics at baseline (date of MRI scan). WHO, World Health Organization; BP, Blood pressure; LV, left ventricle; RV, right ventricle.

Characteristic	n	% or Mean \pm SD
Age (years)		62.9 \pm 14.5
Body surface area (m^2)		1.92 \pm 0.25
Male	169	56
Race		
Caucasian	215	71.2
Asian	7	2.3
Black	13	4.3
Other	28	9.3
Unknown	39	12.9
WHO functional class		
I	1	0
II	45	15
III	214	71
IV	42	14
Haemodynamics		
Systolic BP (mmHg)		131.5 \pm 25.2
Diastolic BP (mmHg)		75 \pm 13
Heart rate (beats/min)		69.8 \pm 22.5
Mean right atrial pressure (mmHg)		9.9 \pm 5.8
Mean pulmonary artery pressure (mmHg)		44.1 \pm 12.6
Pulmonary vascular resistance (Wood units)		8.9 \pm 5.0
Cardiac output (l/min)		4.3 \pm 1.5
LV Volumetry		
LV ejection fraction (%)		61 \pm 11.1
LV end diastolic volume index (ml/m)		110 \pm 37.4
LV end systolic volume index (ml/m)		44 \pm 22.9
RV Volumetry		
RV ejection fraction (%)		38 \pm 13.7
RV end diastolic volume (ml/m)		194 \pm 62
RV end systolic volume (ml/m)		125 \pm 59.3
RV Strain		
Longitudinal (%)		-16.8 \pm 4.7
Radial (%)		+18.0 \pm 4.4
Circumferential (%)		-9.6 \pm 7.0

Table 2

Hyperparameter search ranges for DL network (first column) and optimum hyperparameter values in final model (second column)

Hyperparameter	Search Range	Optimized Value
Dropout	[0.1, 0.9]	0.71
# of nodes in hidden layers	[75, 250]	78
Latent code dimensionality (h)	[5, 20]	13
Reconstruction loss penalty (α)	[0.3, 0.7]	0.6
Learning Rate	$[10^{-6}, 10^{-4.5}]$	$10^{-4.86}$
L_1 regularization penalty	$[10^{-7}, 10^{-4}]$	$10^{-5.65}$