

Generating Healthy Aortic Root Geometries From Ultrasound Images of the Individual Pathological Morphology Using Deep Convolutional Autoencoders

Jannis Hagenah¹, Mohamad Mehdi¹, Floris Ernst¹

¹ Institute for Robotics and Cognitive Systems, University of Luebeck, Germany

Abstract

In valve-sparing aortic root reconstruction surgery, estimating the individual healthy shape of the aortic root as it was before pathological deformation is a challenging task. However, exactly this estimation is necessary to develop personalized aortic root prostheses. To support the surgeon in this decision making, we present a novel approach to reconstruct the healthy shape of an aortic root based on ultrasound images of its pathologically dilated state using representation learning.

The idea is to identify a suitable representation of healthy and pathological aortic root shapes using a supervised variational autoencoder. Then, an image of the dilated root can be encoded, manipulated in the latent space, i.e. shifted towards the distribution of healthy valves, and a synthetic image of this resulting shape can be generated using the decoder.

We evaluate our method on an ex-vivo porcine data set and provide a proof-of-concept of our method in a qualitative and quantitative way. Our results indicate the great potential of reducing a complex shape deformation task to a simple and intuitive shifting towards a specific class. Hence, our method could play an important role in the shaping of personalized implants and is, due to its data-driven nature, not limited to cardiovascular applications but also for other organs.

1. Introduction

Valve-sparing aortic root reconstruction presents a promising treatment for patients suffering from aortic root dilation [1]. This dilation caused by an aneurysm within the aortic root wall dramatically increases the risk of aortic wall rupture and affects the functionality of the aortic valve, leading to valve insufficiency [2]. The idea of the reconstruction surgery is to remove the dilated aortic root in an open surgery and remodel its original, healthy shape using a graft prosthesis. The patient's own aortic valve is spared and sutured into this prosthesis during the surgery.

Even though this technique allows for personalized

treatment, the determination of the individually optimal prosthesis size and shape is challenging. The main reason for this is that only the pathologically dilated state of the aortic root can be assessed while the healthy shape is typically unknown. Hence, an assistive tool supporting the surgeons decision making should be capable of estimating the individual healthy shape of the aortic root only based on images of the pathologically dilated shape. Our goal is the development of such a tool to assist the surgeon in his or her decision making.

As this problem is about inferring from image information, machine learning is a promising approach to solve this problem. However, classical supervised learning techniques, where the prediction model learns the mapping from a dilated shape to the corresponding healthy shape based on annotated shape pairs (comparable to [3]), suffer from the unrealistic demand for training data as collecting these shape pairs in a clinical study is very unlikely. Hence, the prediction model should be trainable on a data set of healthy and dilated aortic roots without the need of knowing both states of one aortic root.

In this work, we present a novel approach to tackle this problem using representation learning. The main idea is to find a good representation of aortic root shapes using a variational autoencoder. Within this representation, i.e. the latent space of the autoencoder, the reconstruction of a healthy shape simplifies to the translation of a dilated sample through the latent space towards the healthy samples. As the decoder network is implicitly trained, it is possible to generate a synthetic medical image of the reconstructed aortic root shape. We introduce this concept of representation learning for planning valve-sparing aortic root reconstructions and present a proof-of-concept on an ex-vivo pig data set.

2. Proposed Concept

The main concept relies on finding a good representation of the aortic root shape. As this representation is learned data-drivenly, it is referred to as the latent space. Each point in this latent space describes one root shape and the pathology changes the shape of the aortic root,

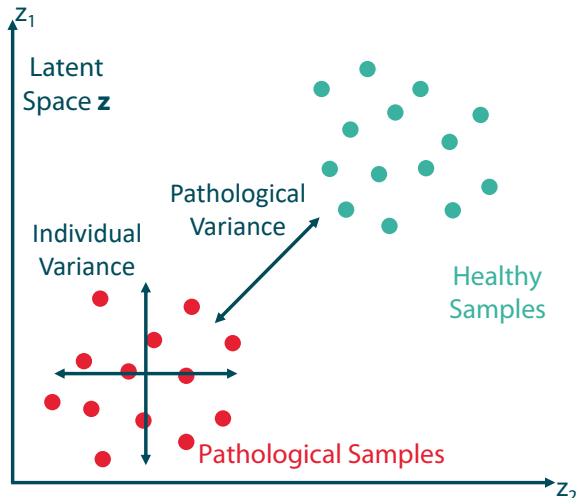


Figure 1. Encoding of pathological and healthy aortic root shapes in a suitable latent space, exemplarily shown in two dimensions.

so all known samples should lie in two distinct clusters, representing healthy and pathological shapes (see Fig. 1). Hence, two different kinds of variance can be observed within the latent space: the individual variance and the pathological variance. The individual variance describes the shape differences from patient to patient and leads to the spreading of the data points within one cluster. The pathological variance describes the change in the anatomy caused by the pathology and describes the difference between the two clusters. If the representation is sufficient, it is clearly visible that a reconstruction of the healthy shape based on the pathological shape can be seen as the translation of a dilated sample along the direction of the pathological variance while sparing as much as possible of the individual variance. Hence, the problem simplifies to a vector translation of a dilated sample through the latent space towards the cluster of healthy shapes (see Fig. 2). One possible translation is the vector that points from the mean of dilated samples to the mean of healthy samples, i.e. from one cluster center to the other one.

Like this, the complex shape deformation in image space can be broken down to a simple vector translation if the representation is sufficient. To find this suitable representation, we propose to use a supervised variational autoencoder. The encoder maps a medical image, for example an ultrasound image, to the latent space and using the decoder, synthetic images of the estimated healthy aortic root shape can be generated.

3. Material & Methods

To validate our proposed concept, we applied the procedure to an ex-vivo porcine data set. We implemented

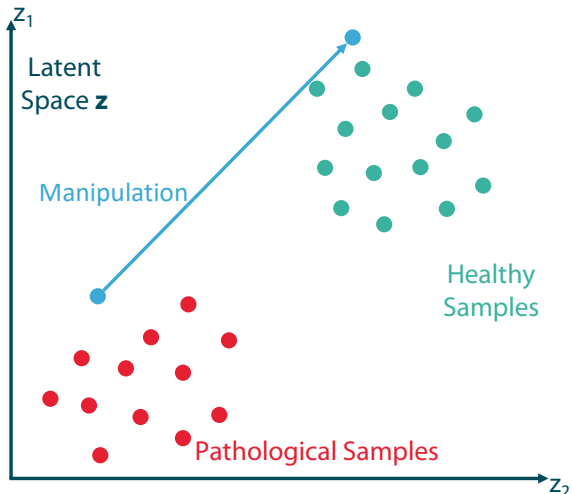


Figure 2. The reconstruction can be performed by translating an encoded pathological sample towards the distribution of healthy root shapes.

the representation learning and evaluated the capability of reconstructing individual healthy shapes, qualitatively and quantitatively. In the following paragraphs, we present the details of this evaluation.

3.1. Data Set

To evaluate our method, we tested it on an ex-vivo porcine data set. The data set was published in [3] and consists of 3D ultrasound images of 24 isolated aortic roots in the closed state under physiologically realistic pressure conditions. During the collection of the data set, at first, images of the natural aortic roots were taken. Then, the roots were manually dilated simulating a pathological aneurysm by making small incisions in the root wall and sewing in additional tissue. Afterwards, images were taken again. Hence, for each of the aortic roots, the healthy as well as the dilated state is known, resulting in 48 ultrasound volumes in total. In this first approach, we extracted one 2D slice image from each volume in which the commissure plane, i.e. the slice image were all three commissure points are visible, is shown [3].

This data set has the big advantage that for each dilated aortic root, the original healthy shape is known. Hence, this ground truth can be addressed during evaluation.

3.2. Network Training

As mentioned above, we propose to utilize a supervised variational autoencoder [4] for finding a sufficient representation. In this study, the encoder network q consists of five convolutional layers with kernel size equals to three [5]. The activation function used within the layers is ReLu

activation. Each convolutional layer is followed by an average pooling layer with a kernel size of two. This enables a smooth transition to the low-dimensional bottleneck representation. The decoder network p is obtained by mirroring the encoder architecture with the average pooling being replaced by upsampling.

The latent space is represented by a dense layer with a size of 16 nodes. As mentioned above, we added a second branch to the network in parallel to the decoder that we call the discriminative branch (d). To this end, we added a dense layer with softmax activation that gets the latent space representation as an input. In an ideal scenario, the output of this branch for an aortic root image should be the correct binary label, i.e. *healthy* or *pathological*.

As a result, we defined a custom loss function that combines classical variational autoencoder loss L_{vae} [4] with the classification loss in the discriminative branch L_d . The loss function L is defined as

$$\begin{aligned} L(w_q, w_p, w_d) &= \alpha L_{vae} + (1 - \alpha) L_d \\ &= \alpha (E(\log p(X|z)) - D_{kl}(q(z|X)||p(z))) \\ &\quad + (1 - \alpha) D_{bce}(d(q(X)), y_{true}), \end{aligned} \quad (1)$$

where X is the input image, y_{true} is the true binary label of this image (*healthy* or *pathological*), w_q , w_p and w_d are the weights of the encoder, decoder and discriminator, respectively, D_{kl} is the Kullback-Leibler-Divergence and D_{bce} is the binary cross entropy [5]. The parameter α controls the weighting between L_{vae} and L_d . In this study, we set $\alpha = 0.5$, so both parts of the loss function were weighted equally.

The implementation and training of the supervised variational autoencoder was done with the help of the *keras* package with *tensorflow* backend [6].

3.3. Reconstruction & Evaluation

In this study, we aimed on a qualitative as well as on a quantitative evaluation. However, quantifying the accuracy of our proposed method is quite hard as a generated image should be compared to a real ultrasound image with sparse structure. In this study, we used the mean squared error (MSE) as well as the structural similarity (SSIM) as metrics to compare pairs of images [7]. To assess the improvement of shape similarity through our reconstruction process, we calculated both error metrics between the pathological and the healthy image as well as between the reconstructed and the healthy image, respectively. As the roots within the images are not aligned and the shape similarity assessment should be rotational and translational in-

variant, we first applied an intensity-based rigid image registration to the image pairs to align the roots before comparing them.

The whole evaluation was performed using a 6-fold cross validation on the roots, i.e. on pairs of pathological and corresponding healthy images. Within the cross validation, the evaluation pipeline was as follows: The representation learning network was trained on the training set and the translation vector t was calculated as the vector pointing from the mean of the pathological roots to the mean of the healthy roots in the latent space description. Then, the pathological images $x_{pathologic}$ of the test set were encoded to their latent space description, manipulated by adding the translation vector t and the resulting point was decoded using the decoder:

$$x_{reco} = p(q(x_{pathologic}) + t) \quad (2)$$

The resulting reconstructed images x_{reco} were compared to their corresponding healthy ground truth as written above.

4. Results

The qualitative analysis revealed that our proposed method was capable of producing realistic aortic root shapes. The reconstructed roots do not show the typical patterns of a pathological dilation and even individual properties of the root geometries are spared during the reconstruction. Fig. 3 exemplarily shows the reconstruction result for three aortic roots. The results of the quantitative analysis are shown in Table 1. By applying our method, the similarity to the healthy image could be greatly increased compared to the pathological image.

Table 1. Quantitative evaluation.

Image Pairs	MSE	SSIM
Pathologic vs. Healthy	502.58	0.76
Reconstructed vs. Healthy	389.0	0.82
Relative Improvement	29.2 %	7.8 %

5. Discussion

The qualitative results indicate the power of the proposed representation learning network as it was capable of reducing the influence of the pathology on the root shape while individual properties were still clearly visible after the reconstruction. We can assume that the latent space description is smooth as all generated images showed realistic aortic root shapes. The quantitative results support these findings. Even though the absolute values of the error metrics are far from optimal, we focus on the improve-




Figure 3. Qualitative reconstruction results for three roots, showing the pathological, the reconstructed and the healthy image.

ment in similarity that we can achieve using our reconstruction as the comparison of ultrasound images with such a sparse structure like an isolated aortic root in a rotational and translational invariant way is very challenging.

Our results indicate that the reconstruction using a quite simple manipulation in a sufficient latent space is possible and produces realistic results that are quite similar to the healthy ground truth. Of course, the artificially induced dilation in the data set does not necessarily mimic a realistic pathology. However, our method is not relying on a handcrafted model but works with variances observed in the data set. Hence, we think that the model should be applicable to real clinical data as well as long as a training set is available. Additional to this step towards clinical application, future work will include a deeper analysis of different architecture types as well as the integration of non-linear manipulations within the latent space. One approach for the latter problem might be to learn these manipulations implicitly from the data set to get rid of a handcrafted manipulation scheme. Another interesting analysis would be the interpretation of the latent space dimensions to understand what kind of variance is encoded in a dimension.

To the best of our knowledge, we presented the first approach to estimate healthy organ shapes based on pathological ones using latent space manipulation. Utilizing representation learning, we were able to reduce a complex shape transformation to a simple and intuitive translation task, opening up a new field of explainable and controllable machine learning methods for surgery planning. We

could achieve a proof-of-concept on a data set of ultrasound images of porcine aortic roots. However, the proposed method is not limited to this application and could be relevant for other challenges within the whole field of personalized prosthetics.

Acknowledgments

The authors would like to thank Dr. Michael Scharfschwerdt from the Department of Cardiac Surgery, University Hospital Schleswig-Holstein for his helpful advise and the support assessing the data set.

References

- [1] Kallenbach K, Karck M, Leyh RG, Hagl C, Walles T, Harringer W, Haverich A. Valve-sparing aortic root reconstruction in patients with significant aortic insufficiency. *The Annals of Thoracic Surgery* 2002;74(5):S1765–S1768.
- [2] Scharfschwerdt M, Sievers HH, Hussein A, Kraatz EG, Mischfeld M. Impact of progressive sinotubular junction dilatation on valve competence of the 3f aortic and sorin solo stentless bioprosthetic heart valves. *European Journal of Cardio Thoracic Surgery* 2010;37(3):631–634.
- [3] Hagenah J, Werrmann E, Scharfschwerdt M, Ernst F, Metzner C. Prediction of individual prosthesis size for valve-sparing aortic root reconstruction based on geometric features. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2016; 3273–3276.
- [4] Kingma DP, Welling M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* 2013;.
- [5] LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 2015; 521(7553):436.
- [6] Chollet F, et al. Keras. <https://keras.io>, 2015.
- [7] Wang Z, Bovik AC, Sheikh HR, Simoncelli EP, et al. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 2004; 13(4):600–612.

Address for correspondence:

Jannis Hagenah
 Institute for Robotics and Cognitive Systems
 University of Luebeck
 Ratzeburger Allee 160, 23562 Luebeck, Germany
 hagenah@rob.uni-luebeck.de