The International Journal of Cardiovascular Imaging (2019) 35:2189–2196
https://doi.org/10.1007/s10554-019-01671-0

ORIGINAL PAPER

Denoising and artefact removal for transthoracic echocardiographic imaging in congenital heart disease: utility of diagnosis specific deep learning algorithms

Gerhard-Paul Diller[1,2,3] · Astrid E. Lammers[1,3,4] · Sonya Babu-Narayan[2] · Wei Li[2] · Robert M. Radke[1] · Helmut Baumgartner[1,3] · Michael A. Gatzoulis[2] · Stefan Orwat[1,3]

Received: 4 April 2019 / Accepted: 14 July 2019 / Published online: 19 July 2019

Abstract

Deep learning (DL) algorithms are increasingly used in cardiac imaging. We aimed to investigate the utility of DL algorithms in de-noising transthoracic echocardiographic images and removing acoustic shadowing artefacts specifically in patients with congenital heart disease (CHD). In addition, the performance of DL algorithms trained on CHD samples was compared to models trained entirely on structurally normal hearts. Deep neural network based autoencoders were built for denoising and removal of acoustic shadowing artefacts based on routine echocardiographic apical 4-chamber views and performance was assessed by visual assessment and quantifying cross entropy. 267 subjects (94 TGA and atrial switch and 39 with ccTGA, 10 Ebstein anomaly, 9 with uncorrected AVSD and 115 normal controls; 56.9% male, age 38.9 ± 15.6 years) with routine transthoracic examinations were included. The autoencoders significantly enhanced image quality across diagnostic subgroups ($p < 0.005$ for all). Models trained on congenital heart samples performed significantly better when exposed to examples from congenital heart disease patients. Our study demonstrates the potential of autoencoders for denoising and artefact removal in patients with congenital heart disease and structurally normal hearts. While models trained entirely on samples from structurally normal hearts perform reasonably in CHD, our data illustrates the value of dedicated image augmentation systems trained specifically on CHD samples.

Keywords Adult congenital heart disease · Congenital heart disease · Deep-learning · Echocardiography · De-noising · Autoencoder

Introduction

Transthoracic echocardiography represents the mainstay in the routine imaging assessment of patients with congenital heart disease (CHD) [1–3]. Especially the longitudinal quantification of ventricular chamber dimensions and function as well as the assessment of valvular function are of interest [2]. With evolving technology, echocardiographic image quality has dramatically improved over the last few decades. While even experienced echocardiographers benefit from this progress, it particularly assists less experienced colleagues in making the correct diagnosis and identifying complications in this setting. Despite these advances image quality is still suboptimal in some patients due to anatomic restrictions or body mass. Recently, the field of machine learning, including deep neuronal (DL) networks, has received increasing attention in medical imaging [4–10]. Frameworks based on DL have been developed for automatic view classification

Gerhard-Paul Diller and Astrid E. Lammers contributed equally to the study.

Electronic supplementary material The online version of this article (https://doi.org/10.1007/s10554-019-01671-0) contains supplementary material, which is available to authorized users.

✉ Gerhard-Paul Diller
gerhard.diller@ukmuenster.de

1 Department of Cardiology III – Adult Congenital and Valvular Heart Disease, University Hospital Muenster, Albert-Schweitzer-Campus 1, 48149 Münster, Germany

2 Adult Congenital Heart Disease Programme, Royal Brompton Hospital, London, UK

3 Competence Network for Congenital Heart Defects, DZHK (German Centre for Cardiovascular Research), Berlin, Germany

4 Division of Paediatric Cardiology, University Hospital Muenster, Münster, Germany

and splitting routine transthoracic investigations into individual views [6]. Automatic chamber segmentation has been described by us and others and could further assist clinicians in evaluating long term changes in ventricular dimensions and function [11, 12]. In addition, DL networks might be employed to augment suboptimal images. We aimed to test the utility of DL based autoencoders in improving echocardiographic image quality in patients with CHD. Specifically, we focused on removing random noise and (artificial) acoustic shadowing artefacts in echocardiographic samples from patients with complex CHD and a biventricular anatomy. As major manufacturers of echocardiographic equipment will likely develop similar tools based on training samples from structurally normal hearts, we also aimed to compare the performance of models trained entirely with structurally normal heart samples to that of models including samples from patients with various types of CHD to assess whether general-purpose algorithms perform well in CHD patients or specific algorithms for CHD would be required for this complex patient population.

Patients and methods

Transthoracic echocardiographic sequences of patients (covering the spectrum of congenital heart disease with a biventricular heart) attending the Adult Congenital Heart Disease Program at the Royal Brompton Hospital London, UK and the Adult Congenital Heart Disease Centre at Muenster University Hospital, Germany with archived recordings were included. In addition, patients with structurally normal hearts and normal ventricular function were identified and served as normal controls. All recordings were performed according to guideline recommendations and loops were stored in a digital DICOM format [13]. Echocardiographic loops were available in an anonymized form and the retrospective use of such de-identified recordings has been approved by the local ethics committees. Before analysis, the echocardiographic apical 4-chamber views were visually assessed for adequate image quality and were automatically converted to a pixel format (bitmap) using custom-written software (MatLab R2018a, MathWorks, Natick, MA).

Simulation of random noise or acoustic shadow artefacts

To simulate random noise, we added a random value sampled from a normal distribution with a mean of zero and a standard deviation of 0.5 to each pixel value in the original image (pixel values in the original image were in the range 0 to 1). Resulting pixel values of <0 or >1 were set to these minimum and maximum values, respectively. For the simulation of acoustic shadow artefacts, random sectors originating close to the presumed original probe position and geometrically aligned with the natural echo-beam orientation were constructed. The apex position of the sector was randomly set within the ± 2 pixels from the horizontal midpoint of the image, and randomly up to 12 pixels vertically away from the presumable probe position. The sector tilt was chosen randomly between $-27°$ and $+27°$ and the width was in the range of 6–14°. All pixel values in the original image within the simulated sector area were set to zero.

Autoencoder architecture

Autoencoders represent a specific form of deep neuronal networks, designed to reconstruct an output from a given input. Internally, they utilize a combination of an encoder, an intermediate (hidden) coding layer and a decoder. In their simplest form autoencoders can be used to reduce the dimensionality of the input data (i.e. for data compression) but they can also be utilized to enhance the output image in various forms. Mathematically, provided a given input x, the autoencoder can be viewed as the combination of an encoder function $g(x)$ and a decoding function $f(g(x))$, reconstructing x based on the results of the encoder function.

A denoising autoencoder is a modification of the autoencoder described above, trained to reconstruct the original image x from a modified version ($\bar{x}$) of the original image x. The modified version of the image $\bar{x}$ could, for example, be obtained by adding random noise to the original image x. Thus, in mathematical notion, this network is represented by $f(g(\bar{x})) \rightarrow x$. When training such a model an appropriate loss function L is required and the model is trained to minimize the difference between the original image (x) and the reconstructed image $f(g(\bar{x}))$, i.e. to minimize $L(x, f(g(\bar{x})))$. Figure 1 illustrates the general architecture of a generic and a denoising autoencoder. The autoencoder setup for removal of sector artefacts is analogous and is shown in Figure A in the online appendix. A standard choice for the loss function in this setting is binary cross-entropy. The concept of cross-entropy has its roots in information theory and it conceptually measures the difference between two probability distributions [14]. For a more theoretical discussion of binary cross-entropy in the setting of autoencoders, we refer to Creswell et al. [15]. Due to its utility in machine learning, in general, this function is implemented in standard neural network libraries such as *keras*. For model comparison, a lower cross-entropy represents a better representation of the underlying distribution and is thus considered more favorable. In addition to cross-entropy, we assessed the quality of image reconstruction based on the sum of squared differences (SSD) between the original image and the reconstructed image. The SSD adds up the squared

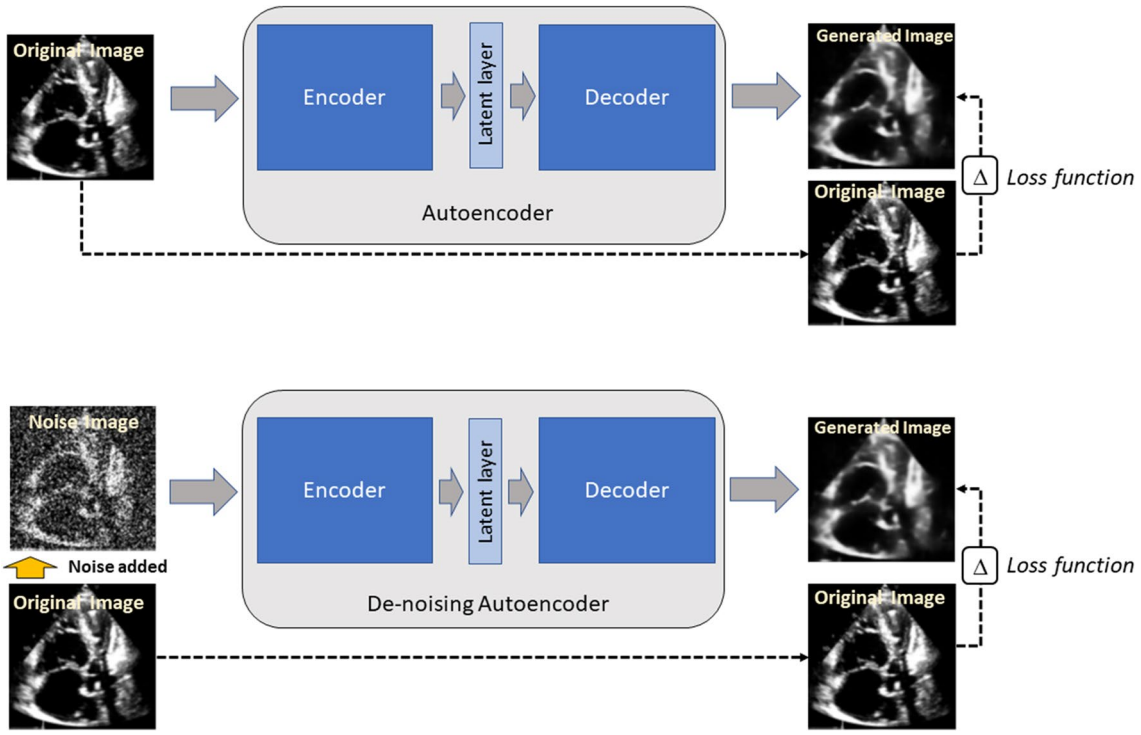

Fig. 1 Upper panel: Illustration of the general architecture of an autoencoder. The input (original) image is processed by an encoder and the resulting latent variable is subsequently decoded back by the decoder sub-network. The generated image is compared with the original image and the model is trained to minimize the difference between the two images using the pre-defined loss function. Lower panel: Illustration of the de-noising autoencoder setup. The original input image is first subjected to the addition of random noise. The resulting noised image is the subsequently presented to the autoencoder setup. The generated image is, however, not compared to the noise image but the original image. Therefore, the model is trained to reconstruct the original image based on the information from the noised image

difference between the brightness value of each pixel in the two respective images. Lower values represent a better reconstruction result.

The realness and anatomical/structural plausibility of the reconstructed images was assessed visually for all frames by two experienced investigators (GPD and SO).

Table 1 illustrates the layer architecture of the autoencoder employed in the current study. Echocardiographic images were down sampled to a resolution of 128×128 pixels with 256 grey-shades. The model was trained over 200–400 epochs with an *adadelta* optimization and binary cross-entropy as a loss function. The batch size was set at 16. The autoencoder networks were constructed in *Tensorflow* using the *keras* package for R [16]. A Intel i7 platform with GPU support (Nvidia GX 1070) was used for training and testing the networks. All analyses were performed using RStudio Version 1.1.456/R-package version 3.5.1.

Table 1 Layer architecture of the autoencoder used

Layer type	Output shape	No. of parameters
Input image	$128 \times 128 \times 1$	0
Conv2D	$128 \times 128 \times 32$	320
MaxPooling2D	$64 \times 64 \times 32$	0
Conv2D	$64 \times 64 \times 32$	9248
MaxPooling2D	$32 \times 32 \times 32$	0
Conv2D	$32 \times 32 \times 32$	9248
UpSampling2D	$64 \times 64 \times 32$	0
Conv2D	$64 \times 64 \times 32$	9248
UpSampling2D	$128 \times 128 \times 32$	0
Conv2D	$128 \times 128 \times 1$	289
Output image	$128 \times 128 \times 1$	0
Total trainable parameters		28,353

Conv2D convolutional 2-D layer, *MaxPooling2D* maximal pooling layer, *UpSampling2D* up-sampling layer

Statistical analysis

Descriptive data are presented as mean ± standard deviation. Comparisons between the performance metrics of different autoencoders were made using a 2-tailed paired Student t-test and a two-sided P-value of < 0.05 was considered indicative of statistical significance.

Results

Overall, we included 152 patients with CHD and a biventricular heart (94 with transposition of the great arteries [TGA] after atrial switch, 39 congenitally corrected TGA, 9 uncorrected atrioventricular septal defects and 10 with Ebstein anomaly) and 115 normal controls. The mean age of the population was 38.9 ± 15.6 years and 56.9% of the subjects were male. For training the models 92,918 apical 4-chamber views of patients with CHD and 20,000 apical 4-chamber view images of subjects with a structurally normal heart were employed. For testing purposes, additional 60,502 images from different patients with CHD and 4354 images from subjects with a structurally normal heart were available. Images used for training and testing the model were from different patients to ensure that the results of the analysis could be generalized to external patients.

Denoising autoencoder

Three different autoencoders were trained based on (i) structurally normal hearts, (ii) samples from CHD patients and (iii) a combination of both, trained firstly on CHD and subsequently of samples from structurally normal hearts. Training these models for 400 epochs required approximately 3 h per model on our local setup. In contrast to training, running the trained autoencoder to reconstruct an individual frame is very fast requiring only approximately 5 ms/frame. Figure 2 illustrates the results of recovering the original image from frames with added white noise for different underlying diagnoses using the three autoencoders described above. As the Figure illustrates, visually, the results look relatively similar and the autoencoder trained entirely on structurally normal hearts provides a reasonable de-noising result in the CHD samples. When quantitatively assessing the results of the de-noising process, however, the two autoencoders trained on CHD samples were significantly superior to the autoencoder trained entirely on normal heart samples as shown in Table 2. Interestingly, even for structurally normal hearts, the models with exposure to CHD hearts during the

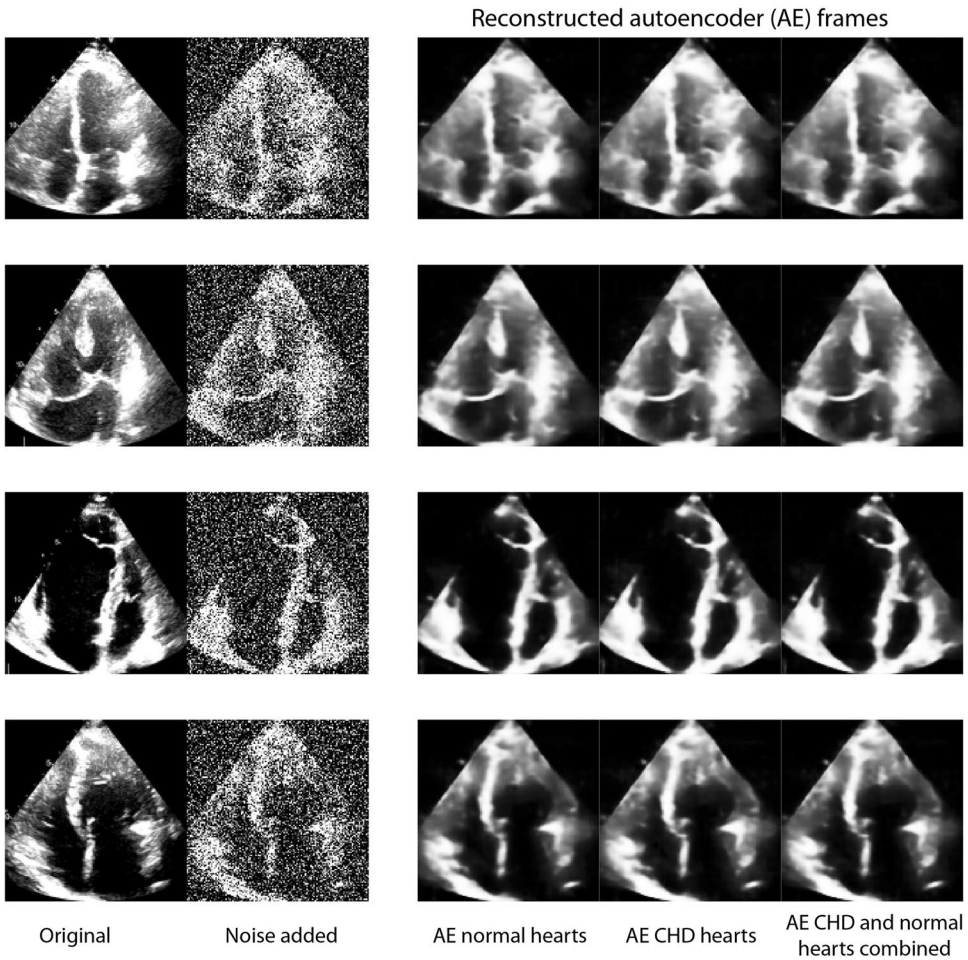

Fig. 2 Representative results of the three autoencoders trained for removing random noise. Example images are from (from top to bottom) from a normal control, a patient with atrioventricular septal defect, an Ebstein anomaly patient and a patient with corrected transposition of the great arteries. *AE* autoencoder. AE trained entirely on normal hearts (=AE normal hearts), AE trained entirely on congenital heart samples (=AE CHD hears) and AE trained jointly on normal and congenital samples (=AE CHD and normal hearts combined)

Table 2 Comparison between different autoencoders (AE) in improving cross-entropy and mean squared difference compared to baseline after adding random white noise

	Cross-entropy			Squared difference		
	Mean ± SD	p value vs. (1)	p value vs. (2)	Mean ± SD	p value vs. (1)	p value vs. (2)
Baseline after adding noise	2.7603 ± 0.3116			1926.86 ± 27.76		
(1) AE trained on normal hearts	0.2912 ± 0.0441			165.93 ± 23.26		
(2) AE trained on CHD hearts	*0.2897 ± 0.0439*	< 0.0001		*160.44 ± 22.96*	< 0.0001	
(3) AE trained on CHD and normal hearts	0.2899 ± 0.0441	< 0.0001	< 0.0001	161.20 ± 23.08	< 0.0001	< 0.0001

Lower values for cross-entropy and squared difference are considered superior. All results are based on the test dataset, originating from different patients than those used for model training. The p-values are based on paired t-test analysis. The best result for each analysis is highlighted by italics

training process performed significantly better compared to the model entirely trained on normal hearts with the model trained jointly on CHD and normal samples providing the best results on normally structured hearts (cross-entropy: 0.2719 ± 0.0321 for the normal heart model, compared to 0.2713 ± 0.0318 for the CHD model and 0.2707 ± 0.0321 for the CHD and normal heart model, $p < 0.0001$ for all comparisons; mean square difference: 151.97 ± 13.59 for the normal heart model, compared to 149.06 ± 13.60 for the CHD and 147.43 ± 13.25 for the combined CHD and normal heart model, $p < 0.0001$ for all comparisons). The process of denoising can also be applied to an entire loop as illustrated by the movie file in the online only material.

Shadow artefact autoencoders

To assess the utility of autoencoders to remove acoustic drop-out shadow artefacts we trained three different autoencoders as mentioned above. Again, one model was trained entirely on samples from structurally normal hearts, one on samples from CHD patients and one on a combination of both, being trained firstly on CHD and subsequently of samples from structurally normal hearts. Training these models for 200 epochs required approximately 90 min per model on our local setup. As above, running the trained autoencoder to reconstruct an individual frame required 6 ms/frame. Figure 3 illustrates the results of recovering the original image from frames with an artificially added sector to simulate drop-out artefacts for different underlying diagnoses using the three autoencoders described above. The results look relatively similar on visual assessment, however, quantitatively the two autoencoders trained on CHD samples were again significantly superior to the autoencoder trained entirely on normal heart samples as shown in Table 3.

When assessing the model performance on a test set of 4354 samples from normal subjects not used for model development, the two autoencoder models trained on CHD samples (entirely or in part), performed significantly better compared to the model developed entirely on normal samples (cross entropy: normal heart autoencoder 0.2649 ± 0.0369 vs. 0.2597 ± 0.0327 for the CHD model and 0.2587 ± 0.0328 for the combined CHD and structurally normal heart model, $p < 0.0001$ for all comparisons; mean squared difference: normal heart autoencoder 133.89 ± 79.06, vs. 118.86 ± 61.52 for the CHD model and 113.49 ± 60.54 for the combined CHD/structurally normal model, $p < 0.0001$ for all comparisons, respectively).

Discussion

The current study demonstrates the feasibility of automatic denoising and removal of segmental acoustic artefacts using novel deep learning-based algorithms. Autoencoders trained on a large number of frames from patients with CHD and structurally normal hearts significantly improved image quality across the spectrum of CHD in patients with biventricular hearts. When comparing different autoencoders, we found that training autoencoders on samples from CHD and normal subjects significantly improved results. This significant improvement was also evident in samples from structurally normal hearts, suggesting that developing such algorithms based on data sets including CHD samples may be beneficial even outside the setting of CHD.

Life-long surveillance is recommended for most CHD patients with medium and high complexity underlying lesions and transthoracic echocardiography remains the cornerstone of longitudinal assessment in this setting [1, 3]. Due to complex anatomy, associated extracardiac lesions, concomitant lung disease and occasionally abnormal situs and location of the heart in the thoracic cavity obtaining adequate image quality may be challenging even for experienced echocardiographers [2]. Using technological solutions to improve image quality is appealing and may, in selected patients, help to avoid expensive and less widely available investigations such as cardiac magnetic resonance imaging. The latter is also often not available in the setting of patients with an implanted cardiac pacemaker or defibrillator. Based

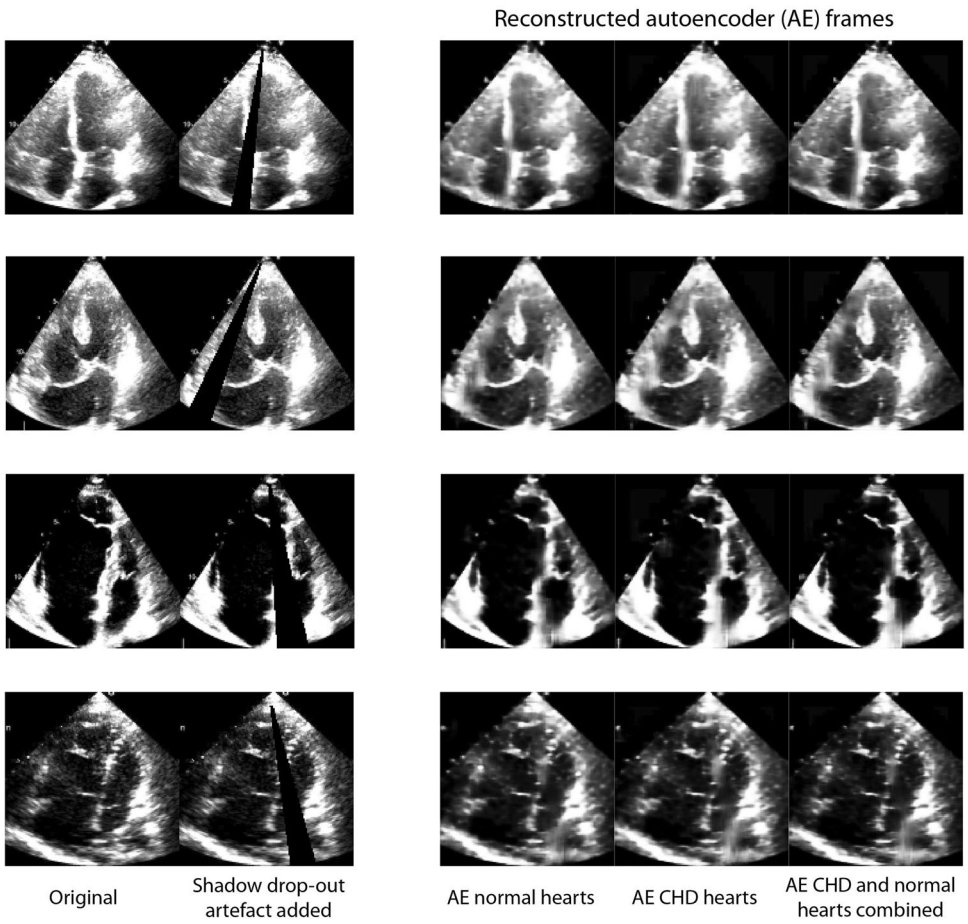

Fig. 3 Representative results of the three autoencoders trained for removing shadowing artefacts. Example images are from (from top to bottom) from a normal control, a patient with atrioventricular septal defect, an Ebstein anomaly patient and a patient with atrial switch operation for transposition of the great arteries. *AE* autoencoder. AE trained entirely on normal hearts (=AE normal hearts), AE trained entirely on congenital heart samples (=AE CHD hears) and AE trained jointly on normal and congenital samples (=AE CHD and normal hearts combined)

Table 3 Comparison between different autoencoders (AE) in improving cross-entropy and mean squared difference compared to baseline after adding sector drop-out artefacts

	Cross-entropy			Squared difference		
	Mean ± SD	p value vs. (1)	p value vs. (2)	Mean ± SD	p value vs. (1)	p value vs. (2)
Baseline after adding noise	1.5518 ± 0.5904			453.29 ± 242.22		
(1) AE trained on normal hearts	0.2809 ± 0.0463			133.59 ± 73.89		
(2) AE trained on CHD hearts	0.2754 ± 0.0434	< 0.0001		113.31 ± 50.97	< 0.0001	
(3) AE trained on CHD and normal hearts	*0.2744 ± 0.0435*	< 0.0001	< 0.0001	*107.00 ± 49.87*	< 0.0001	< 0.0001

Lower values for cross-entropy and squared difference are considered superior. All results are based on the test dataset, originating from different patients than those used for model training. The p-values are based on paired t-test analysis. The best result for each analysis is highlighted by italics

on deeper physical understanding and using improved technology such as tissue harmonic imaging echocardiographic image quality has dramatically improved in the last decades [17]. The current study illustrates how deep learning algorithms may assist in further augmenting image quality in the setting of congenital heart disease. Given the rapid progress in this area, it is to be expected that even more impressive results may be obtainable using these frameworks in the years to come. We believe that this technology has great potential and may assist non-congenital echocardiographers in obtaining the correct diagnosis. Beyond the potential for assisting human operators, increasingly, automated systems for view classification, diagnosis identification and segmentation based on routine echocardiograms are being developed. Madani et al. have presented a deep-learning based solution for automatic view classification of echocardiograms, splitting the investigation into separate views and thus assisting with analysis and measurement

[6]. In addition, we have recently described a methodology for automatic classification of the underlying diagnosis and automatic segmentation of the systemic right ventricle in subjects with normal hearts and patients with systemic right ventricles. The accuracy of the segmentation process was comparable to experienced human investigators in our previous paper [11]. These neuronal networks are generally trained on echocardiographic frames of adequate quality and we contend that analyzing data with reduced image quality may be challenging. Combining de-noising autoencoders as presented in the current study upfront with these segmenting networks may alleviate the problem and allow networks not specifically trained on noisy images to perform adequately. Abdi et al. have described a framework to quantify the quality of transthoracic echocardiograms using convolutional networks, a system that may also be incorporated in the routine echocardiographic workflow in future [18].

The current study is—to our knowledge—the first to assess the utility of deep learning algorithm based autoencoders in patients with congenital heart disease. The results of this study, however, have implications beyond congenital heart disease as models (co-)trained on samples from patients with structurally abnormal hearts performed significantly better even in the context of structurally normal hearts. We speculate that by incorporating a wider anatomical spectrum, models are exposed to a more heterogeneous input and thus learn a more abstract representation of the sample space, eventually improving results. Heterogeneity is generally useful in the setting of training neuronal networks and random geometric transformations are frequently applied to improve model performance in this setting [12]. The methods described here may also have the potential to improve image quality and remove artefacts in cardiac magnetic resonance imaging. In fact, motion suppression and the removal of spacio-temporal artefacts in congenital heart disease using deep neuronal networks has been reported [19]. Overall, it is likely that postprocessing of echocardiographic images will gain importance and carries the promise of efficiency gains and improved diagnostic accuracy.

Limitations

We included echocardiographic images obtained from two tertiary centers for congenital heart disease. Therefore, we cannot exclude the possibility that image quality in a community setting maybe lower and this may affect autoencoder performance. We limited our analyses to patients with biventricular hearts and to apical 4-chamber views to evaluate the potential utility of autoencoders in patients with CHD. Further studies, including other echocardiographic views and patients with univentricular hearts should be performed to extend on these findings.

We limited our analyses to single 2-D frames. Theoretically, it would be desirable to include a full stack of images covering an entire heart cycle into the analysis to capitalize on existing cross-correlations between frames. This, however, requires further experience with autoencoders in this setting and more performant computers, to apply the methods described here directly to a 4-dimensional data set. Segmental drop-out artefacts, by definition, represent a local loss of information in the image frame and this information cannot be fully recovered by any technological solution. The results of the segmental autoencoder, therefore, should be considered best guess estimates of the original image. The main application of this technology is likely as an upfront addition to automatic segmentation networks, likely to be affected by such artefacts. While we present highly significant differences between autoencoders it remains currently unclear what absolute differences in cross-entropy and SSD would be of clinical relevance.

Before implementation of the algorithms presented here into routine clinically used commercial packages, these tools should undergo further testing, including accepted phantom models to assess performance in a real-world setting.

Conclusions

The current study demonstrates for the first time the utility of deep learning frameworks for denoising and artefact removal in patients with congenital heart disease. While models trained entirely on samples from structurally normal hearts perform reasonably in congenital heart disease, our data illustrates the value of dedicated image augmentation systems trained specifically on congenital samples. In addition, our data suggest that incorporating a sufficient proportion of congenital samples in the training process also improves the performance of autoencoders in the setting of structurally normal hearts.

Acknowledgement This study was supported by a research Grant from the EMAH Stiftung Karla Voellm, Krefeld, Germany and by the German Competence Network for Congenital Heart Defects (Funded by the German Federal Ministry of Education and Research, BMBF -FKZ 01G10210, 01GI0601).

Author contribution GPD, AEL and SO planned and conducted the study. GPD, AEL and SO prepared and analyzed the data using DL networks. SBN, RR, WL, MG and HB made substantial contributions in analysis, drafting the article and revising it critically for important intellectual content. All authors gave final approval of the version to be submitted and any revised version.

Funding This study was supported by a research Grant from the EMAH Stiftung Karla Voellm, Krefeld, Germany. The Adult Congenital Heart Centre and Centre for Pulmonary Hypertension, Royal Brompton Hospital, London, UK have received support from Actelion

UK, Pfizer UK, GSK UK, the British Heart Foundation and the NIHR Cardiovascular and Respiratory Biomedical Research Units.

Compliance with ethical standards

Conflict of interest The authors have no conflicts of interest.

References

1. Baumgartner H, Bonhoeffer P, De Groot NM, de Haan F, Deanfield JE, Galie N, Gatzoulis MA, Gohlke-Baerwolf C, Kaemmerer H, Kilner P, Meijboom F, Mulder BJ, Oechslin E, Oliver JM, Serraf A, Szatmari A, Thaulow E, Vouhe PR, Walma E, Task Force on the Management of Grown-up Congenital Heart Disease of the European Society of C, Association for European Paediatric C, Guidelines ESCCfP (2010) ESC guidelines for the management of grown-up congenital heart disease (new version 2010). Eur Heart J 31(23):2915–2957
2. Li W, West C, McGhie J, van den Bosch AE, Babu-Narayan SV, Meijboom F, Mongeon FP, Khairy P, Kimball TR, Beauchesne LM, Ammash NM, Veldtman GR, Oechslin E, Gatzoulis MA, Webb G (2018) Consensus recommendations for echocardiography in adults with congenital heart defects from the International Society of Adult Congenital Heart Disease (ISACHD). Int J Cardiol 272:77–83
3. Stout KK, Daniels CJ, Aboulhosn JA, Bozkurt B, Broberg CS, Colman JM, Crumb SR, Dearani JA, Fuller S, Gurvitz M, Khairy P, Landzberg MJ, Saidi A, Valente AM, Van Hare GF (2018) 2018 AHA/ACC guideline for the management of adults with congenital heart disease. Circulation 139:e698–e800
4. Johnson KW, Torres Soto J, Glicksberg BS, Shameer K, Miotto R, Ali M, Ashley E, Dudley JT (2018) Artificial intelligence in cardiology. J Am Coll Cardiol 71(23):2668–2679
5. LeCun Y, Bengio Y, Hinton G (2015) Deep learning. Nature 521(7553):436–444
6. Madani A, Arnaout R, Mofrad M, Arnaout R. Fast and accurate view classification of echocardiograms using deep learning. npj Digit Med 2018;**1**(1):6.
7. Madani A, Ong JR, Tibrewal A, Mofrad MR. Deep echocardiography: data-efficient supervised and semi-supervised deep learning towards automated diagnosis of cardiac disease. npj Digit Med 2018;**1**(1):59.
8. Rajan SP, Kavitha V (2017) Diagnosis of cardiovascular diseases using retinal images through vessel segmentation graph. Curr Med Imaging Rev 13(4):454–459
9. Ronneberger O, Fischer P, Brox T (2015) U-net: convolutional networks for biomedical image segmentation. In: International conference on medical image computing and computer-assisted intervention. Springer, New York, pp 234–241
10. Salem A-BM, Revett K, El-Dahshan E-SA (2009) Machine learning in electrocardiogram diagnosis. In: International multiconference on computer science and information technology, 2009. IMCSIT'09. IEEE, pp 429–433
11. Diller GP, Babu-Narayan S, Li W, Radojevic J, Kempny A, Uebing A, Dimopoulos K, Baumgartner H, Gatzoulis M, Orwat S (2019) Utility of machine learning algorithms in assessing patients with a systemic right ventricle. Eur Heart J Cardiovasc Imaging. https://doi.org/10.1093/ehjci/jey211
12. Bai W, Sinclair M, Tarroni G, Oktay O, Rajchl M, Vaillant G, Lee AM, Aung N, Lukaschuk E, Sanghvi MM, Zemrak F, Fung K, Paiva JM, Carapella V, Kim YJ, Suzuki H, Kainz B, Matthews PM, Petersen SE, Piechnik SK, Neubauer S, Glocker B, Rueckert D (2018) Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. J Cardiovasc Magn Reson 20(1):65
13. Lang RM, Bierig M, Devereux RB, Flachskampf FA, Foster E, Pellikka PA, Picard MH, Roman MJ, Seward J, Shanewise JS, Solomon SD, Spencer KT, Sutton MS, Stewart WJ, Chamber Quantification Writing G, American Society of Echocardiography's G, Standards C, European Association of E (2005) Recommendations for chamber quantification: a report from the American Society of Echocardiography's Guidelines and Standards Committee and the Chamber Quantification Writing Group, developed in conjunction with the European Association of Echocardiography, a branch of the European Society of Cardiology. J Am Soc Echocardiogr 18(12):1440–1463
14. Rubinstein R (1999) The cross-entropy method for combinatorial and continuous optimization. Methodol Comput Appl Probab 1(2):127–190
15. Creswell A, Arulkumaran K, Bharath AA (2017) On denoising autoencoders trained to minimise binary cross-entropy. arXiv:1708.08487
16. Chollet F (2015) Keras. GitHub.
17. Turner SP, Monaghan MJ (2006) Tissue harmonic imaging for standard left ventricular measurements: fundamentally flawed? Eur J Echocardiogr 7(1):9–15
18. Abdi AH, Luong C, Tsang T, Allan G, Nouranian S, Jue J, Hawley D, Fleming S, Gin K, Swift J, Rohling R, Abolmaesumi P (2017) Automatic quality assessment of echocardiograms using convolutional neural networks: feasibility on the apical four-chamber view. IEEE Trans Med Imaging 36(6):1221–1230
19. Hauptmann A, Arridge S, Lucka F, Muthurangu V, Steeden JA (2018) Real-time cardiovascular MR with spatio-temporal artifact suppression using deep learning—proof of concept in congenital heart disease. Magn Reson Med 81:1143–1156