Minimising Multi-Centre Radiomics Variability Through Image Normalisation: A Pilot Study

Víctor M. Campello (✉ victor.campello@ub.edu)

University of Barcelona

Carlos Martín-Isla

University of Barcelona

Cristian Izquierdo

University of Barcelona

Andrea Guala

Vall d'Hebron Institut de Recerca (VHIR)

José F. Rodríguez-Palomares

Vall d'Hebron Institut de Recerca (VHIR)

David Viladés

Hospital de la Santa Creu i Sant Pau, Universitat Autònoma de Barcelona

Martín L. Descalzo

Hospital de la Santa Creu i Sant Pau, Universitat Autònoma de Barcelona

Mahir Karakas

University Heart & Vascular Center Hamburg

Ersin Çavus

University Heart & Vascular Center Hamburg

Zahra Raisi-Estabragh

Barts Heart Centre, St Bartholomew's Hospital, Barts Health NHS Trust

Steffen E. Petersen

Barts Heart Centre, St Bartholomew's Hospital, Barts Health NHS Trust

Sergio Escalera

University of Barcelona

Santi Seguí

University of Barcelona

Karim Lekadir

University of Barcelona

Research Article

Keywords:

Posted Date: January 27th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1250808/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Minimising multi-centre radiomics variability through image normalisation: A pilot study

Víctor M. Campello^{1,*}, Carlos Martín-Isla¹, Cristian Izquierdo¹, Andrea Guala^{2,3}, José F. Rodríguez Palomares^{2,3,4}, David Viladés⁵, Martín L. Descalzo⁵, Mahir Karakas^{6,7}, Ersin Çavuş^{6,7}, Zahra Raisi-Estabragh^{8,9}, Steffen E. Petersen^{8,9,10,11}, Sergio Escalera^{1,12}, Santi Seguí¹, and Karim Lekadir¹

¹ *Artificial Intelligence in Medicine Lab (BCN-AIM), Barcelona, Spain*

² *Vall d'Hebron Institut de Recerca (VHIR), Barcelona, Spain*

³ *CIBER-CV, Instituto de Salud Carlos III, Madrid, Spain*

⁴ *Department of Cardiology, Hospital Universitari Vall d'Hebron, Barcelona, Spain*

⁵ *Cardiac Imaging Unit, Cardiology Service, Hospital de la Santa Creu i Sant Pau, Universitat Autònoma de Barcelona, Barcelona, Spain*

⁶ *Dept. of Cardiology, University Heart & Vascular Center Hamburg, Hamburg, Germany*

⁷ *DZHK (German Center for Cardiovascular Research), Germany*

⁸ *Barts Heart Centre, St Bartholomew's Hospital, Barts Health NHS Trust, London, United Kingdom*

⁹ *William Harvey Research Institute, NIHR Barts Biomedical Research Centre, Queen Mary University London, London, United Kingdom*

¹⁰ *Health Data Research UK, London, United Kingdom*

¹¹ *Alan Turing Institute, London, United Kingdom*

¹² *Computer Vision Center, Universitat Autònoma de Barcelona, Barcelona, Spain*

* *Corresponding author: victor.campello@ub.edu*

ABSTRACT

Radiomics is an emerging technique for the quantification of imaging data that has recently shown great promise for deeper phenotyping of cardiovascular disease. Thus far, the technique has been mostly applied in single-centre studies. However, one of the main difficulties in multi-centre imaging studies is the inherent variability of image characteristics due to centre differences. In this paper, a comprehensive analysis of radiomics variability under several image- and feature-based normalisation techniques was conducted using a multi-centre cardiovascular magnetic resonance dataset. 218 subjects divided into healthy (n=112) and hypertrophic cardiomyopathy (n=106, HCM) groups from five different centres were considered. First and second order texture radiomic features were extracted from three regions of interest, namely the left and right ventricular cavities and the left ventricular myocardium. Two methods were used to assess features' variability. First, feature distributions were compared across centres to obtain a distribution similarity index. Second, two classification tasks were proposed to assess: 1) the amount of centre-related information encoded in normalised features (centre identification) and 2) the generalisation ability for a classification model when trained on these features (healthy versus HCM classification). The results showed that the feature-based harmonisation technique ComBat is able to remove the variability introduced by centre information from radiomic features, at the expense of degrading classification performance. Piecewise linear histogram matching normalisation gave features with greater generalisation ability for classification (0.75 ± 0.09). Models trained with features from images without normalisation showed the worst performance overall (0.47 ± 0.13). In conclusion, centre-related information removal did not imply good generalisation ability for classification.

Introduction

For the last decade, there has been a great amount of research devoted to identifying and improving quantitative image biomarkers for precise diagnosis, risk assessment and patient stratification for different pathologies. In particular, radiomics seems to be a promising technique to quantify image-derived biomarkers based on shape, intensity and higher-order texture patterns for a region of interest defined a priori, since it is able to characterise image patterns that are hardly visible to the naked eye.

These computer-extracted features have the potential to perform an exhaustive analysis of medical images as shown in the literature, predominantly in oncology¹ but also more recently for neurodevelopmental disorders² or cardiovascular disease³. However, radiomic features have proven to be highly sensitive to changes in scanning protocols and scanner manufacturers, resulting in a limited reproducibility^{4,5} (see also the exhaustive reviews by Yip and Aerts⁶ and by Traverso et al.⁷) and thus posing an important problem that needs to be solved before implementing these techniques in clinical practice. Despite this, the majority of previous research considered single-institution datasets, due in part to the difficulty in obtaining imaging studies from multiple centres. More recently, several works using multi-centre studies have assessed the robustness of this technique (see for example Raisi-Estabragh et al.⁸, for a test-retest study). A number of works have proposed harmonisation guidelines for computed tomography (CT) or positron emission tomography in multi-centre scenarios, while no guideline is available for magnetic resonance imaging (MRI), where the lack of a standard intensity grayscale – such as Hounsfield units in CT – poses further difficulty (see Da Ano et al.⁹ and references therein).

All previous multi-centre MRI radiomics studies focused either on brain or cancer imaging. Due to the lack of multi-centre cardiac imaging radiomics literature, a detailed introduction about brain and cancer imaging is presented. Two types of techniques are used to standardise features across institutions: image- and feature-based transformations.

At the image level, the most common techniques are image intensity normalisation (mean subtraction and division by the standard deviation) or image intensity rescaling to a fixed range (usually from 0 to 1). Other more sophisticated techniques exist, such as bias field correction, isotropic resampling, histogram matching and piecewise linear histogram matching (PLHM). Finally, some techniques are inherently defined for brain imaging and were not considered in this study. Um et al. used T1-weighted MRI brain scans to assess radiomics variability across two different institutions after five image preprocessing techniques were applied, including global and region of interest (ROI) rescaling, bias correction, isotropic resampling and histogram matching⁵. They concluded that histogram matching is the best technique for reducing feature variability and successfully discriminate between different patient subgroups with glioblastoma. Isaksson et al. evaluated the effect of four normalisation techniques on classification performance to identify prostate cancer in T2-weighted MRI¹⁰. The normalisation method that resulted in the best classification accuracy was the PLHM transformation using intensities from healthy prostate as reference instead of the whole image to extract landmarks. Finally, Carré et al. standardised brain MRI studies using three different intensity normalisation techniques in order to find their effects on radiomics robustness, being image intensity normalisation the technique that yielded the best results¹¹.

At the feature level, Chatterjee et al. improved the robustness of radiomics from images of primary uterine adenocarcinoma by applying feature normalisation for each institution dataset independently¹². Orlhac et al., instead, used the empirical Bayes harmonisation method – also referred to as ComBat¹³ – to remove inter-centre variability¹⁴. The transformed features resulted in a sensitivity increase for distinguishing between Gleason grades in prostate cancer studies and in similar distributions for features from brain scans for 1.5T and 3T machines.

In this work, a multi-centre cardiac MRI dataset was considered to analyse the effect of several image- and feature-based normalisation techniques over radiomic features variability and model generalisation across institutions.

Results

Feature variability

Texture features variability showed a great disparity depending on the preprocessing method under consideration (see Fig. 1). For both the end-diastole (ED) and the end-systole (ES) frames, the highest percentage of features with similar distributions across institutions (Jensen-Shannon divergence, JSD, below 0.01) was obtained when applying ROI-based PLHM or ROI-based rescaling as shown in Figure 1 (in between 62% and 76% of features). In contrast, the proportion of features below the 0.01 JSD threshold was the lowest for the same methods, but applied at the whole image level, and for images without any normalisation (HM, PLHM and O in Fig. 1). The proportion of features below the given threshold was reduced to less than 54%, indicating that feature distributions are less similar in these cases. Additionally, the large standard deviation, represented by the black horizontal bars in Figure 1, was associated to differences depending on the ROIs and, especially, on the centre pairs being compared (see Fig. S1 and S2 in the supplementary material for a detailed comparison of these factors).

The application of the ComBat harmonisation method had an averaging effect, reducing the proportion of features with close distributions for the most robust methods in the previous paragraph and increasing it for the least robust methods, as shown in Figure 2. Thus, smaller differences were found between methods after the application of ComBat.

Notably, in both cases (with and without ComBat), the feature variability was similar for first and second order features for most methods.

Figure 1. Percentage of first and second order features below the 0.01 JSD threshold distinguished for healthy subjects. Results are averaged over centre pairs and ROI and presented separately for ED and ES frames. The black lines represent the standard deviation. O: original images (without normalisation), R: image intensity recaling, N: image intensity normalisation, HM: histogram matching and PLHM: piecewise linear histogram matching. An “R.” in front of a method means that it is applied at ROI level.

Among all feature families, Gray Level Size Zone Matrix (GLSZM), Gray Level Run Length Matrix (GLRLM) and Gray Level Dependence Matrix (GLDM) presented the highest dissimilarities among distributions after the application of ROI-based PLHM normalisation in general, as demonstrated by the greater JSD values in Table 1. Gray Level Co-occurrence Matrix (GLCM) and first order features obtained the best similarity scores. As noted above, the JSD was averaged over all family of features to an approximate value of 0.011 after the application of ComBat. The features found with the most dissimilar distributions in both cardiac time frames, ED and ES, prior to the application of ComBat were zone variance, large area emphasis and large area low gray level emphasis (GLSZM), kurtosis (1st order), large dependence low gray level emphasis

Figure 2. Percentage of first and second order features below the 0.01 JSD threshold distinguished for healthy subjects after the application of the feature-based harmonisation tool ComBat. Results are averaged over centre pairs and ROI and presented separately for ED and ES frames. The black lines represent the standard deviation. O: original images (without normalisation), R: image intensity recaling, N: image intensity normalisation, HM: histogram matching and PLHM: piecewise linear histogram matching. An “R.” in front of a method means that it is applied at ROI level.

(GLDM) and long run low gray level emphasis (GLRLM).

Table 1. Mean and standard deviation (in parenthesis) for JSD for distributions of features obtained after the application of R.PLHM normalisation on healthy patients. Results are presented separately for ED and ES frames and for each feature family before and after the application of ComBat harmonisation. Values are averaged over ROI.

Family	Without Combat		With Combat	
	ED	ES	ED	ES
1st order	0.008 (0.007)	0.008 (0.007)	0.012 (0.012)	0.011 (0.008)
GLCM	0.008 (0.006)	0.008 (0.007)	0.011 (0.011)	0.012 (0.011)
GLDM	0.010 (0.009)	0.009 (0.008)	0.012 (0.012)	0.011 (0.009)
GLRLM	0.010 (0.010)	0.009 (0.007)	0.010 (0.009)	0.011 (0.010)
GLSZM	0.012 (0.012)	0.011 (0.010)	0.012 (0.011)	0.011 (0.009)

R.PLHM: ROI-based piecewise linear histogram matching; ED: End-diastole; ES: End-systole.

Centre identification

When assessing the centre information encoded in the extracted features, second order texture features carried more information in general than first order features, as demonstrated by the accuracy differences for classifiers built on them in Figure 3 (orange and blue boxes). Features from original images (without normalisation) were the most discriminative features with testing accuracy above 0.88 (± 0.03 –0.07) for the three ROIs under consideration and for both feature types, first order and texture features. When comparing normalisation techniques at whole image level, the centre information was preserved almost equally for all methods and all heart ROIs (Fig. 3, orange and blue boxes in the top row).

When normalisation was applied at ROI level, larger differences appeared depending on the method and the ROI under

consideration (Fig. 3, orange and blue boxes in the bottom row). ROI-based PLHM consistently reduced the ability of models to infer the centre of origin for each sample of features, achieving the lowest performance in most cases for features without ComBat harmonisation. Finally, ComBat harmonisation was able to remove centre information from features almost entirely for most normalisation techniques and original images, as shown by the red and green boxes in Figure 3 around 0.2 (the accuracy obtained by a random model over five centres).

Figure 3. Accuracy of random forest models when predicting the centre of origin for first and second order texture features before and after the application of ComBat harmonisation. The row above corresponds to image preprocessing techniques applied at the whole image level, while in the row below they are applied at the ROI level. O: original images (without normalisation), R: image intensity recaling, N: image intensity normalisation, HM: histogram matching and PLHM: piecewise linear histogram matching. An “R.” in front of a method means that it is applied at ROI level.

Generalisation

With regards to the patient classification task into healthy and hypertrophic cardiomyopathy (HCM) groups in unseen centres, the best accuracy was obtained when using features extracted after applying the PLHM transformation and the worst accuracy was associated to features from original images (Fig. 4, upper row). However, models trained for ROI-based normalisation techniques showed small variability in performance between them (Fig. 4, bottom row).

In more detail, for models trained with features from Vall d’Hebron studies, the highest accuracy was 0.745 (median: 0.778 [0.694, 0.826]), obtained after PLHM without the application of ComBat. When ComBat harmonisation was used, the highest accuracy was obtained after the application of the same image normalisation technique but was reduced to 0.697 (median: 0.702 [0.630, 0.759]). For models trained with features from Sagrada Familia studies, the best accuracies were again obtained for PLHM and were 0.767 (median: 0.787 [0.696, 0.833]) and 0.763 (median: 0.766 [0.739, 0.796]) before and after the application of ComBat harmonisation, respectively. For these models, features mostly from the myocardium (MYO) were

among the most important features for the model prediction according to the Gini importance¹⁵. The top 20 most important features contained mean and median intensity, kurtosis and skewness (1st order), joint average and autocorrelation (GLCM) and run length non-uniformity and long run high gray level emphasis (GLRLM).

When comparing the accuracy between validation (same institution) and testing (unseen institutions) sets, models that obtained the highest accuracy on validation did not generalise well to new unseen centres (Fig. 5). Importantly, models trained with features from ROI-based normalisation methods showed poorer generalisation than methods at the whole image level. Within normalisation methods at the whole image level, features extracted after PLHM obtained the best testing accuracy despite their lower performance in validation when compared to other techniques (median: 0.800 [0.78, 0.90] when trained with Vall d'Hebron and median: 0.71 [0.67, 0.93] when trained with data from Sagrada Familia).

Figure 4. Accuracy of random forest models on unseen centres for classification of HCM versus healthy patients. All models were trained with a combination of first and second order texture features from all ROIs. The first column corresponds to models trained with features extracted from Vall d'Hebron studies, while models in the second column were trained with features from Sagrada Familia studies. The row above corresponds to image preprocessing techniques applied at the whole image level, while in the row below they are applied at the ROI level. HCM: Hypertrophic cardiomyopathy, O: original images (without normalisation), R: image intensity recaling, N: image intensity normalisation, HM: histogram matching and PLHM: piecewise linear histogram matching. An “R.” in front of a method means that it is applied at ROI level.

Discussion

Radiomic features are promising biomarkers for better disease characterisation. However, their variability across centres makes it difficult to establish reproducible biomarkers based on them⁶. In this study, a comprehensive analysis was carried out to assess feature variability across centres as well as model generalisation for a classification task after the application of several image normalisation techniques and a feature-based harmonisation technique (ComBat).

Based on the results presented, ROI-based PLHM is a good normalisation technique to preserve similar feature distributions across domains (see Fig. 3) and to reduce the amount of centre-related information encoded in radiomic features compared to original images (see Fig. 1 and 2). In brain MRI literature, however, the transformations that yielded less feature variability and more similar distributions were histogram matching and image intensity normalisation^{5,11}, although Um et al. did not consider PLHM in their work.

At feature normalisation level, ComBat satisfactorily removed centre-related information to the point that models were not able to discriminate between features depending on the institution of origin of each scan. As a drawback, the final feature distributions at different centres were less similar than before the application of ComBat according to the JSD.

Figure 5. Accuracy of random forest models on the validation set (same domain) versus the testing set (unseen centres) for classification of HCM versus healthy patients. Results are presented without ComBat harmonisation. All models were trained with a combination of first and second order texture features from all ROIs. The first column corresponds to models trained with features extracted from Vall d'Hebron studies, while models in the second column were trained with features from Sagrada Familia studies. The row above corresponds to image preprocessing techniques applied at the whole image level, while in the row below they are applied at the ROI level. HCM: Hypertrophic cardiomyopathy, O: original images (without normalisation), R: image intensity recaling, N: image intensity normalisation, HM: histogram matching and PLHM: piecewise linear histogram matching. An “R.” in front of a method means that it is applied at ROI level.

Interestingly, features with large area/dependence/run and low gray level emphasis were among the features with greater variability across centres. These features share the same definition and emphasise the importance of large regions with low intensities. Hence, according to the results, these large dark regions are not meaningful across centres. Lastly, healthy and HCM groups showed disparities in the results, being the second group the one that presented less similar distributions. This could be explained by the diversity in hypertrophy severity across cohorts.

With regards to generalisation ability, models trained with features from original images resulted in poor performance for differentiation of healthy subjects from HCM patients in unseen centres, highlighting the importance of normalisation

techniques for multi-centre studies. Models trained with features extracted after PLHM normalisation obtained the best accuracy. In this method, average landmarks are obtained for a reference population, while for histogram matching the reference was only one subject. This could explain why these methods showed differences in performance despite relying on the same principle, since defining a template using only one subject could introduce unwanted bias in the analysis. Moreover, the selection of a particular population or subject as reference template in these methods may affect the results, especially for histogram matching (see Fig. S3 and S4 in the supplementary material).

Importantly, successful centre-related information removal from radiomic features does not imply greater generalisation ability. In fact, ROI-based PLHM and ComBat harmonisation methods were not among the best generalisation techniques for the HCM classification task (Fig. 4, bottom row). When compared with the brain MRI literature, Orlhac et al. did find an improvement in sensitivity for differentiating between low and high risk patient groups when using ComBat harmonisation, although the authors did not compare different image normalisation techniques¹⁴. Lastly, the model trained after PLHM, which showed the best generalisation ability, obtained medium performance on the validation set signalling a reduction in overfitting. The most important features for this model were predominantly features from the MYO, which made sense for the classification task at hand, since HCM is most evident when looking at the myocardium.

This work presents several limitations. First, the dataset was not perfectly balanced across the five centres and the population was not controlled by age, sex, body size, or myocardial volume, which could result in dissimilarities across feature distributions.

The choice of HCM classification as a metric for generalisation has some drawbacks, since the heart suffers morphological changes and some texture features are known to be correlated with the shape¹⁶. This could contribute to overestimating the generalisation ability. The inclusion of other pathologies that greatly affect the myocardium, such as myocarditis or infarction, would potentially result in a less biased generalisation loss estimation.

Finally, according to the Imaging Biomarker Standardization Initiative, ISBI, second order features from different texture matrices may be modelled better with different intensity discretisation levels (*e.g.* GLSZM are better characterised for low discretisation levels while it is the opposite for GLCM)¹⁷. In this work, the same discretisation level was used for all features.

Conclusions

In summary, this study showed that centre-related information removal does not imply good generalisation ability for classification. ComBat harmonisation was able to remove centre-related information from radiomic features satisfactorily, while showing limited generalisation ability. PLHM normalisation resulted in the best generalisable model for classification of healthy subjects from HCM patients. The choice of reference template when performing histogram matching may affect the results. PLHM was robust against a change of reference population. Finally, the radiomic features from GLSZM, GLDM and GLRLM families showed greater variability than first order and GLCM features. Further studies with a larger sample size are needed in order to replicate the results presented and to assess the effect of different biological covariates.

Material and methods

Data and feature extraction

For this work, a subset of the Multi-Centre, Multi-Vendor & Multi-Disease Cardiac Image Segmentation Challenge (M&Ms) dataset was considered¹⁸. These studies were collected in the context of the M&Ms Challenge and all patients signed the informed consent. The study protocol was approved by the Ethical Committee for Clinical Research for each institution involved, and it follows the ethical guidelines of the Declaration of Helsinki. Since the dataset is publicly available, no direct approval from the ethical committee was necessary on our side. In particular, 218 cardiac magnetic resonance studies were selected from the five available centres accounting for healthy subjects as well as patients with hypertrophic cardiomyopathy (HCM). The exact distribution across the five centres is presented in Table 2. All the scanners considered had a field strength of 1.5T and the averaged in-plane resolution ranged from 0.85 to 1.45 millimetres. More detailed information about the scanners used can be found in Table 3.

Table 2. Distribution of diseases per centre considered in the analysis

	Creu Blanca	Dexeus	Sagrada Familia	Universitätsklinikum Hamburg-Eppendorf	Vall d'Hebron	Total
Vendor	Canon	General Electric	Philips	Philips	Siemens	
Healthy	14	11	33	32	22	112
HCM	15	5	37	14	25	106

HCM: hypertrophic cardiomyopathy

Table 3. Average specifications for the studies acquired in the five different centres.

Centre	Vendor	Model	In-plane resolution (mm)	Slice thickness (mm)	Number of slices	Intesities range
Vall d'Hebron	Siemens	Magnetom Avanto	1.32	9.2	12	0 – 1193
Sagrada Familia	Philips	Achieva	1.20	9.9	10	0 – 357
Universitätsklinikum Hamburg-Eppendorf	Philips	Achieva	1.45	9.9	11	0 – 3725
Dexeus	General Electric	Signa Excite	1.36	10	12	0 – 3030
Creu Blanca	Canon	Vantage Orian	0.85	10	13	0 – 14442

Each study consisted of a short-axis cine cardiac magnetic resonance volume. Segmentations of three anatomical ROIs, the left and right ventricle cavities (LV and RV, respectively) and the left ventricle myocardium (MYO), were provided for two temporal phases, ES and ED. The delineations were revised to follow the same Standard Operating Procedure to avoid inter-centre variability. Consequently, the induced bias on radiomic features by inter-centre variability in delineations was negligible.

Radiomic features were extracted using the PyRadiomics library¹⁹, version 3.0.1. Prior to the extraction, all images were resized to match the same spatial resolution of $1 \times 1 \text{ mm}^2$, since radiomic features have been shown to intrinsically depend on voxel size and on the number of voxels¹⁶. Fixed bin widths of 25 and 0.05 were used during feature extraction for images before and after normalisation, respectively. This resulted in a good balance between number of bins and computing requirements. The number of bins after normalisation ranged between 20 and 80, depending on the intensity values for each ROI. Only images without normalisation gave a large variability in terms of number of bins (from 14 to 570).

A total of 100 features were extracted per ROI. They include shape features and first and second order texture features. In this work, only texture features were used, since shape depends only on the ROI segmentation and not on the image intensity. First order texture features refer to commonly used statistical metrics to describe the histogram of intensity values such as mean, minimum, maximum, kurtosis, skewness, entropy and energy, among others. Second order texture features are statistical measures extracted from the four texture matrices considered in this library: Gray Level Co-occurrence Matrix (GLCM), Gray Level Size Zone Matrix (GLSZM), Gray Level Run Length Matrix (GLRLM) and Gray Level Dependence Matrix (GLDM). These features account for different details in the spatial coarseness, variability, heterogeneity and symmetry of textures. A complete list of the features considered is included in the supplementary material.

Normalisation techniques

Four normalisation techniques were considered at the image level:

- R: image intensity rescaling to the range 0 – 1,
- N: image intensity normalisation (mean subtraction and division by the standard deviation),
- HM: histogram matching using scikit-image²⁰, version 0.17.2,
- PLHM: piecewise linear histogram matching²¹, also referred to as Nyúl-Udupa normalisation.

For the histogram matching transformation, an image intensity histogram is interpolated so that it matches a template histogram. In this work, a subject was selected visually from Sagrada Familia as the template after ensuring that the image did not present artifacts. For the PLHM transformation, the code implementation by²² was employed. In this case, a batch of images from one centre was needed to obtain the averaged histogram deciles (landmarks) that were then used as reference for the transformation of new image histograms. The landmarks were computed for studies from Sagrada Familia. All transformations were applied both to the whole image and at ROI level, independently. Data from Sagrada Familia were used as reference, since it was the centre with the greatest number of scans.

Regarding feature-based normalisation techniques, the empirical Bayes harmonisation method proposed by Johnson and Rabinovic (ComBat) was considered^{13,23}. This method assumes that the contributions to the final feature values can be separated in biological covariates (*e.g.* pathology) and centre effects (*e.g.* different scan manufacturers). Then, the empirical Bayes method is used to estimate the distributions for these terms from original data and adjust the final feature values to remove centre effects. The ComBat method is robust against outliers and does not need large sample sizes for each centre batch, what makes it a good option for the current study. However, feature distributions are assumed to follow normal distributions for each centre separately, a requirement not always satisfied by the data. For this reason, a quantile transformation (scikit-learn²⁴, version 0.23.2) had to be applied to all radiomic features for each institution independently before ComBat could be used (we used 20 as number of quantiles). No covariates were used along with ComBat harmonisation.

Variability assessment

Radiomics variability across centres was assessed by computing the Jensen-Shannon divergence (JSD) between pairs of feature distributions obtained for healthy subjects within the different ROIs. The HCM pathological group was not considered in this analysis since the possible existence of different HCM sub-groups could introduce uncontrolled bias to the results. The JSD gives a positive measure of how similar two distributions are, with 0 the value obtained when the two distributions are identical. A threshold of 0.01 JSD was selected based on the median of the overall distribution as the relative point where changes in feature proportions were to be assessed.

Then, in order to analyse model generalisation, two tasks were proposed. First, the amount of centre-encoded information after the application of each normalisation technique was measured by training Random Forest (RF) models to identify the source centre for each feature set. The hypothesis was that features with less information about their centre of origin should be more difficult to differentiate and thus, more similar between centres, enhancing the generalisation. Secondly, model generalisation was assessed directly by training RF models for patient classification into healthy or HCM groups, for each normalisation technique. RF were chosen over other techniques due to their simplicity to train and their effectiveness to model non-linear relations between input and output. For all cases, a random seed was fixed before training each model to make the results reproducible.

For the centre identification task, models were trained either with first order or second order features as input variables for each ROI separately (LV, MYO and RV). A five-fold cross-validation was used for obtaining an estimate of the average classification accuracy with reduced bias. Notably, in this case, a lower accuracy represents that features carry less centre information.

For the patient classification task, models were trained with a combination of first and second order features from all three ROIs, so that RF was able to select the most predictive features during training. Five runs of the same cross-validation scheme were considered in this case (five different random seeds), since the variability in models was higher and the accuracy estimates showed greater bias. In particular, these models were trained with features from only one dataset (Vall d'Hebron ($n = 38$) or Sagrada Familia $n = 56$, as these centres had a greater number of samples) and tested on the other four. No feature selection was conducted prior to model training. The most important features for the best performing models were obtained with the mean impurity decrease method¹⁵, also called Gini importance. For these models, a greater accuracy represents better generalisation.

Competing interests statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. SEP provides consultancy to and is shareholder of Circle Cardiovascular Imaging, Inc., Calgary, Alberta, Canada.

Author contributions statement

VMC, CMI, CI, SE, SS and KL contributed to the conceptualization of this work. VMC, CMI, CI and SS designed the methodology. VMC, CMI and CI developed the software tools necessary for preprocessing and analysing image files. VMC, CMI, CI, AG, JFRP, DV, MLD, MK and EÇ participated in the data collection and curation process. All authors reviewed the manuscript.

Funding

This work was partly funded by the European Union's Horizon 2020 research and innovation program under grant agreement number 825903 (euCanSHare project). KL received funding from the Spanish Ministry of Science, Innovation and Universities under grant agreement RTI2018-099898-B-I00. SE received funding by the Spanish project PID2019-105093GB-I00 (MINECO/FEDER, UE) and ICREA under the ICREA Academia programme. AG has received funding from the Spanish Ministry of Science, Innovation and Universities (IJC2018-037349-I). ZRE was supported by a British Heart Foundation Clinical Research Training Fellowship (FS/17/81/33318). SEP acknowledges support from the National Institute for Health Research (NIHR) Biomedical Research Centre at Barts. SEP acknowledges support from the "SmartHeart" EPSRC programme grant (www.nihr.ac.uk; EP/P001009/1). SEP also acknowledges support from the CAP-AI programme, funded by the European Regional Development Fund and Barts Charity.

Acknowledgments

The authors would like to thank Violeta Pina for useful discussions about the contents of this manuscript.

Data Availability Statement

The datasets analysed in this study can be found openly at the M&Ms Challenge webpage: ub.edu/mnms.

References

1. Aerts, H. J. W. L. *et al.* Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nat. Commun.* **5**, DOI: [10.1038/ncomms5006](https://doi.org/10.1038/ncomms5006) (2014).
2. Sun, H. *et al.* Psychoradiologic utility of MR imaging for diagnosis of attention deficit hyperactivity disorder: A radiomics analysis. *Radiology* **287**, 620–630, DOI: [10.1148/radiol.2017170226](https://doi.org/10.1148/radiol.2017170226) (2018).

3. Cetin, I. *et al.* A radiomics approach to computer-aided diagnosis with cardiac cine-mri. In *Statistical Atlases and Computational Models of the Heart. ACDC and MMWHS Challenges*, 82–90, DOI: [10.1007/978-3-319-75541-0_9](https://doi.org/10.1007/978-3-319-75541-0_9) (Springer International Publishing, Cham, 2018).
4. Lee, J. *et al.* Radiomics feature robustness as measured using an MRI phantom. *Sci. Reports* **11**, DOI: [10.1038/s41598-021-83593-3](https://doi.org/10.1038/s41598-021-83593-3) (2021).
5. Um, H. *et al.* Impact of image preprocessing on the scanner dependence of multi-parametric MRI radiomic features and covariate shift in multi-institutional glioblastoma datasets. *Phys. Medicine & Biol.* **64**, 165011, DOI: [10.1088/1361-6560/ab2f44](https://doi.org/10.1088/1361-6560/ab2f44) (2019).
6. Yip, S. S. F. & Aerts, H. J. W. L. Applications and limitations of radiomics. *Phys. Medicine Biol.* **61**, R150–R166, DOI: [10.1088/0031-9155/61/13/r150](https://doi.org/10.1088/0031-9155/61/13/r150) (2016).
7. Traverso, A., Wee, L., Dekker, A. & Gillies, R. Repeatability and reproducibility of radiomic features: A systematic review. *Int. J. Radiat. Oncol.* **102**, 1143–1158, DOI: [10.1016/j.ijrobp.2018.05.053](https://doi.org/10.1016/j.ijrobp.2018.05.053) (2018). Imaging in Radiation Oncology.
8. Raisi-Estabragh, Z. *et al.* Repeatability of cardiac magnetic resonance radiomics: A multi-centre multi-vendor test-retest study. *Front. Cardiovasc. Medicine* **7**, DOI: [10.3389/fcvm.2020.586236](https://doi.org/10.3389/fcvm.2020.586236) (2020).
9. Da-Ano, R., Visvikis, D. & Hatt, M. Harmonization strategies for multicenter radiomics investigations. *Phys. Medicine & Biol.* **65**, 24TR02, DOI: [10.1088/1361-6560/aba798](https://doi.org/10.1088/1361-6560/aba798) (2020).
10. Isaksson, L. J. *et al.* Effects of MRI image normalization techniques in prostate cancer radiomics. *Phys. Medica* **71**, 7–13, DOI: [10.1016/j.ejmp.2020.02.007](https://doi.org/10.1016/j.ejmp.2020.02.007) (2020).
11. Carré, A. *et al.* Standardization of brain MR images across machines and protocols: bridging the gap for MRI-based radiomics. *Sci. Reports* **10**, DOI: [10.1038/s41598-020-69298-z](https://doi.org/10.1038/s41598-020-69298-z) (2020).
12. Chatterjee, A. *et al.* Creating robust predictive radiomic models for data from independent institutions using normalization. *IEEE Transactions on Radiat. Plasma Med. Sci.* **3**, 210–215, DOI: [10.1109/trpms.2019.2893860](https://doi.org/10.1109/trpms.2019.2893860) (2019).
13. Johnson, W. E., Li, C. & Rabinovic, A. Adjusting batch effects in microarray expression data using empirical bayes methods. *Biostatistics* **8**, 118–127, DOI: [10.1093/biostatistics/kxj037](https://doi.org/10.1093/biostatistics/kxj037) (2006).
14. Orlhac, F. *et al.* How can we combat multicenter variability in MR radiomics? validation of a correction procedure. *Eur. Radiol.* DOI: [10.1007/s00330-020-07284-9](https://doi.org/10.1007/s00330-020-07284-9) (2020).
15. Breiman, L. Random forests. *Mach. Learn.* **45**, 5–32, DOI: [10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324) (2001).
16. ul Hassan, M. S. *et al.* Voxel size and gray level normalization of CT radiomic features in lung cancer. *Sci. Reports* **8**, DOI: [10.1038/s41598-018-28895-9](https://doi.org/10.1038/s41598-018-28895-9) (2018).
17. Zwanenburg, A. *et al.* The image biomarker standardization initiative: Standardized quantitative radiomics for high-throughput image-based phenotyping. *Radiology* **295**, 328–338, DOI: [10.1148/radiol.2020191145](https://doi.org/10.1148/radiol.2020191145) (2020).
18. Campello, V. M. *et al.* Multi-centre, multi-vendor and multi-disease cardiac segmentation: The M&Ms challenge. *IEEE Transactions on Med. Imaging* DOI: [10.1109/TMI.2021.3090082](https://doi.org/10.1109/TMI.2021.3090082) (2021).
19. van Griethuysen, J. J. *et al.* Computational radiomics system to decode the radiographic phenotype. *Cancer Res.* **77**, e104–e107, DOI: [10.1158/0008-5472.can-17-0339](https://doi.org/10.1158/0008-5472.can-17-0339) (2017).
20. van der Walt, S. *et al.* scikit-image: image processing in python. *PeerJ* **2**, e453, DOI: [10.7717/peerj.453](https://doi.org/10.7717/peerj.453) (2014).
21. Nyúl, L. G. & Udupa, J. K. On standardizing the MR image intensity scale. *Magn. Reson. Medicine* **42**, 1072–1081, DOI: [10.1002/\(sici\)1522-2594\(199912\)42:6<1072::aid-mrm11>3.0.co;2-m](https://doi.org/10.1002/(sici)1522-2594(199912)42:6<1072::aid-mrm11>3.0.co;2-m) (1999).

22. Reinhold, J. C., Dewey, B. E., Carass, A. & Prince, J. L. Evaluating the impact of intensity normalization on MR image synthesis. In Angelini, E. D. & Landman, B. A. (eds.) *Medical Imaging 2019: Image Processing*, DOI: [10.1117/12.2513089](https://doi.org/10.1117/12.2513089) (SPIE, 2019).
23. Fortin, J.-P., Sweeney, E. M., Muschelli, J., Crainiceanu, C. M. & Shinohara, R. T. Removing inter-subject technical variability in magnetic resonance imaging studies. *NeuroImage* **132**, 198–212, DOI: [10.1016/j.neuroimage.2016.02.036](https://doi.org/10.1016/j.neuroimage.2016.02.036) (2016).
24. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [suppmaterial.pdf](#)