



HAL
open science

Characterizing myocardial ischemia and reperfusion patterns with hierarchical manifold learning

Benoit Freiche, Patrick Clarysse, Magalie Viallon, P. Croisille, Nicolas Duchateau

► **To cite this version:**

Benoit Freiche, Patrick Clarysse, Magalie Viallon, P. Croisille, Nicolas Duchateau. Characterizing myocardial ischemia and reperfusion patterns with hierarchical manifold learning. Statistical Atlases and Computational Modeling of the Heart (STACOM), MICCAI Workshop, 2021, Strasbourg, France. hal-03510191

HAL Id: hal-03510191

<https://hal.science/hal-03510191>

Submitted on 4 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Characterizing myocardial ischemia and reperfusion patterns with hierarchical manifold learning

Benoît Freiche¹, Patrick Clarysse¹, Magalie Viallon^{1,2}, Pierre Croisille^{1,2}, and Nicolas Duchateau¹

¹ Univ Lyon, Université Claude Bernard Lyon 1, INSA-Lyon, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69621, Lyon, France

² Department of Radiology, Hôpital Nord, University Hospital of Saint-Étienne, France

Abstract. We aim at better understanding the mechanisms of ischemia and reperfusion, in the context of acute myocardial infarction. For this purpose, imaging and in particular magnetic resonance imaging are of great value in the clinic, but the richness of the images is currently underexploited. In this paper, we propose to characterize myocardial ischemia and reperfusion patterns across a population beyond the scalar measurements used in the clinic. Specifically, we adapted representation learning techniques to not only characterize the population distribution in terms of scar and microvascular obstruction patterns, but also regarding the appearance of late gadolinium images which reflects tissue heterogeneity. To do so, we implemented a hierarchical manifold learning approach where the embedding from a higher-level content (the images) is guided by one from a lower-level content (the infarct and microvascular obstruction segmentations). We demonstrate its relevance on 1711 late gadolinium enhancement slices from 123 patients with acute ST-elevation myocardial infarction. We designed ways to balance the contribution of each level in the hierarchy, and quantify its impact on the overall distribution and on sample neighborhoods. We notably observe that the obtained latent space is a balanced contribution between the two levels of the hierarchy, and is more robust to challenging images subjected to artifacts or specific lesion patterns.

Keywords: Manifold learning · Decision Hierarchy · Information fusion · Cardiac imaging · Myocardial infarction

1 Introduction

The ischemic mechanisms following the obstruction of a coronary artery can lead to both structural and functional myocardial damage. For acute myocardial infarction, the benefits of treatments that restore the coronary circulation are counterbalanced by potential reperfusion injuries (microvascular obstruction, MVO) due to a sudden blood reflow in areas that were deprived of it [2]. In this context,

imaging and in particular cardiac magnetic resonance imaging plays a crucial role to understand the mechanisms of ischemia-reperfusion [5]. However, due to limited analysis tools, the richness of the acquired images is under-exploited in the clinical analysis. The lesion characteristics are limited to simple scalar descriptors (extent and transmuralty, mainly) [1], and the image contents are not exploited (e.g. the heterogeneity of pixel values within the segmented lesions).

The field of representation learning offers very efficient tools to better characterize the lesion patterns within a population. It allows mapping the high-dimensional data (e.g. images) to a simplified latent representation that facilitates the analysis of individual or subgroup trends. Within this field, manifold learning offers a sound framework where statistical distances in the latent space can be exploited for such a mapping. It assumes that the input samples lie on a (non-linear) mathematical manifold that is unknown but can be learnt from data.

However, the analysis of the imaging content within the myocardium is not straightforward. For instance, in late gadolinium enhancement (LGE) images, MVO is displayed as dark areas within bright and larger regions indicating the infarct. The gray levels of MVO and healthy tissues may be close (see Fig. 1), which can fool the metric used to compare images (often in a pixel-wise fashion), in particular for large MVO. Other critical issues can happen in case of image artifacts. This analysis could be substantially more robust by using additional imaging information, up to the local appearance of the acquired images.

To merge the information from different imaging descriptors, several fusion strategies exist within the field of manifold learning. For instance, Multiple Manifold Learning (MKL) [8] finds the best linear combination between the affinity matrices associated to each descriptor. A non-linear fusion process called Similarity Network Fusion (SNF) [11] has also been proposed to iteratively merge several descriptors. Nonetheless, these methods perform the early fusion of all the descriptors at the same time, which may be suboptimal in our context. A better integration scheme could consist in a hierarchical learning process, to guide the embedding from a given descriptor by a previous embedding from a lower-level descriptor. This approach is already part of standard clinical reasoning through decision trees [10] and of computational methods through random forests [7]. However, it is hardly scalable to multiple high-dimensional descriptors from images. *Bhatia et al.* [4] proposed an interesting hierarchical manifold learning scheme that could overcome this data integration problem, but it has only been exploited to study a single medical imaging modality at different resolutions. In contrast, we want to incorporate the data from several medical imaging modalities in a hierarchical way and estimate a single representation for a population of patients.

In this paper, we aim at improving the analysis of tissue heterogeneity in LGE images by prior knowledge from the segmented lesions, which can be seen as a hierarchical way of estimating a latent space. We propose to use the hierarchical manifold learning framework such that the embedding from a higher level content (the LGE images) is guided by the one from a lower level content (the infarct and

microvascular obstruction zones segmented on these images). Here, we keep the problem unsupervised and examine the distribution of ischemia and reperfusion patterns across a population. As there is no groundtruth for such an unsupervised problem, we design ways of selecting relevant hyperparameters a posteriori. We demonstrate its relevance on a population of 123 patients with acute myocardial infarct, to improve the analysis of potentially subtle tissue heterogeneities beyond the prior segmentation masks.

2 Method

2.1 Hierarchical representation problem statement

We build a two-level hierarchical model, where \mathbf{x}_i^0 stands for the i -th sample from the lower/parent level (the image of the segmented lesions, where pixel values lie in the interval $[0,2]$, with 0, 1 and 2 respectively standing for the healthy myocardium, the infarct, and the MVO), and \mathbf{x}_i^1 corresponds to the same sample from the higher level (the LGE image). We aim at estimating the higher/child level latent space $\mathbf{Y}^1 = [\mathbf{y}_i^1]_{i \in [1,K]}$ guided by the lower level latent space $\mathbf{Y}^0 = [\mathbf{y}_i^0]_{i \in [1,K]}$, K being the number of samples.

2.2 Spectral embedding

In this work, manifold learning is achieved within the diffusion maps framework [6], which served as a basis for fusion [11, 8] and hierarchical [4] algorithms.

For each level $m = \{0, 1\}$ in the hierarchy, pairwise affinities between individuals are encoded within the matrix $\mathbf{W}^m = [W_{ij}^m] \in \mathbb{R}^{K \times K}$ defined as:

$$W_{ij}^m = \begin{cases} \exp(-\frac{\|\mathbf{x}_i^m - \mathbf{x}_j^m\|^2}{2\sigma^2}) & \text{if } j \in \mathcal{N}_k(i), \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where $\mathcal{N}_k(i)$ stands for the neighborhood of the i -th sample, based on the k closest samples. The graph Laplacian is defined from this matrix as $\mathbf{L}^m = \mathbf{D}^m - \mathbf{W}^m$, where \mathbf{D}^m is a diagonal matrix such that $D_{ii}^m = \sum_j W_{ij}^m$.

Diffusion maps consists in performing the spectral decomposition of the graph Laplacian to estimate \mathbf{Y}^m . In practice, this is achieved by diagonalizing the matrix $\tilde{\mathbf{W}}^m = (\mathbf{D}^m)^{-\frac{1}{2}} \mathbf{W}^m (\mathbf{D}^m)^{-\frac{1}{2}}$, which corresponds to working with the normalized graph Laplacian. It can be seen as a Markov chain matrix, where \tilde{W}_{ij}^m represents the probability of moving from sample i to j in one step of a random walk on the graph [6]. The embedding corresponds to the eigenvectors associated to the first higher eigenvalues of $\tilde{\mathbf{W}}^m$, after removing the trivial case associated to the eigenvalue 1. It stands for the principal directions of diffusion across the data manifold, approximated by the graph made of the available samples.

2.3 Hierarchical Manifold Learning

The hierarchical embedding proposed in *Bhatia et al.* [4] builds upon this framework and minimizes the following cost function:

$$\arg \min_{\mathbf{Y}^1} (1 - \mu) \sum_i \sum_j \|\mathbf{y}_i^1 - \mathbf{y}_j^1\|^2 W_{ij}^1 + \mu \sum_i \|\mathbf{y}_i^1 - \mathbf{y}_i^0\|^2, \quad (2)$$

where $\mathbf{Y}^0 = [\mathbf{y}_i^0]$ was previously estimated by applying diffusion maps to the lower-level data, and $\mu \in [0, 1]$ balances the contributions of the higher and lower levels to the hierarchical embedding. If $\mu = 1$, the optimal solution is $\mathbf{Y}^1 = \mathbf{Y}^0$, so the hierarchical embedding is in this case the lower-level embedding. If $\mu = 0$, it corresponds to performing diffusion maps on the higher level only. In their paper, *Bhatia et al.* showed that there is an analytic solution to this cost function minimization for $\mu \neq 0$:

$$\mathbf{Y}^1 = (\mu \mathbf{I} + 2(1 - \mu) \mathbf{L}^1)^{-1} \mu \mathbf{Y}^0, \quad (3)$$

where \mathbf{I} stands for the identity matrix.

2.4 Hyperparameters optimization

Bhatia et al. arbitrarily set the weighting parameter μ . In contrast, we propose two complementary strategies to find the best embedding for our application.

First, we computed a-posteriori each term in the energy function (Eq. 2), and defined the optimal μ as the value for which the two terms are balanced (Fig. 2b).

In addition, we quantified the point-to-point distances between the estimated hierarchical embedding \mathbf{Y}^1 and the embeddings estimated for the higher and lower levels considered independently. To reduce bias in the distances, we rescaled the embeddings globally so that the standard deviations along their first dimension match, and determined the sign of the eigenvectors that produced the best match. The optimal μ using this second strategy corresponds to the embedding at equal distance from the high and low levels (Fig. 2c).

We implemented the method in Python 3.7.6. The whole algorithm was computed on an standard laptop within a few seconds. The limiting part is the computation of the affinity matrices for the whole set of images. The hierarchical part of the algorithm (Eq. 2) is really quick as it doesn't require any optimization step.

3 Experiments and Results

3.1 Data and preprocessing

We analyzed the data from 123 patients with acute myocardial infarction recruited in the MIMI study [3]. The myocardial content of the LGE images was resampled to a reference anatomy using atlas-based techniques, as done in [9].

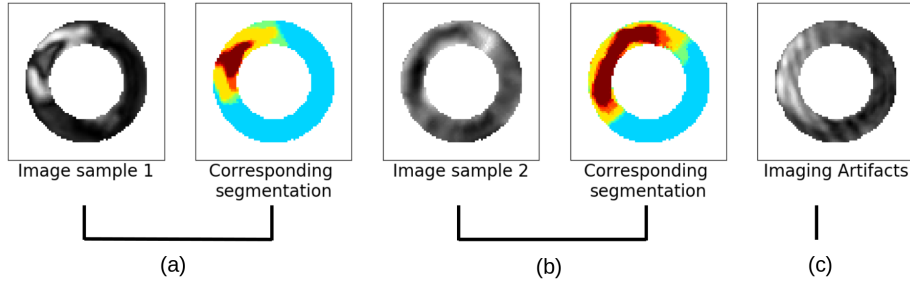


Fig. 1: Representative pairs of LGE images and their corresponding segmentation. The MVO and the infarct respectively correspond to the dark/red and the bright/yellow regions in the grayscale/segmented images. Some challenging samples are displayed: a very large MVO that covers most of the lesion and could be confounded with healthy myocardium (b), and a slice with typical MRI artifacts (c). Note that during the transport of imaging data to a common reference, imaging contents may have been interpolated within a slice and across slices, leading to non-categorical values for the segmented images (not completely blue/yellow/red).

In brief, this consists in a pixel-wise parameterization based on the local radial, circumferential, and long-axis coordinates, which only requires labelling the LV-RV junction on each slice, and the identification of the apex and base levels in the stack of slices.

In this paper, we analyzed the 2D slices independently to benefit from a larger dataset. Besides, the image stacks were artificially reoriented for patients with LCX or RCA infarcts to match the location of the LAD infarct subgroup and therefore prevent the infarct location from confounding the analysis. As a result, we analyzed 1711 samples of segmented images as the parent, and their corresponding grayscale images as the child (see representative examples in Fig. 1).

3.2 Latent space organization

We first applied diffusion maps to the segmented images in which the infarct and MVO are labeled, leading to the lower-level embedding \mathbf{Y}^0 (Fig. 2a-b-c, left column). Then, we computed the affinity matrix associated to the LGE image contents to define the graph Laplacian \mathbf{L}^1 , and estimated the higher-level embedding \mathbf{Y}^1 from Eq. 3 for several values of μ spanning the interval $[0, 1]$. For comparison purposes, diffusion maps were also directly applied to the LGE images (Fig. 2a-b-c, right column).

The bandwidth σ of the kernel involved in the affinity matrices \mathbf{W}^m was set experimentally for each latent space. In the literature, it is commonly set as the average distance between a sample and its k -th nearest neighbor. However, this choice was not relevant in our case, in particular with MVO patterns. For instance, if σ is too small, the principal directions of diffusion can be strongly

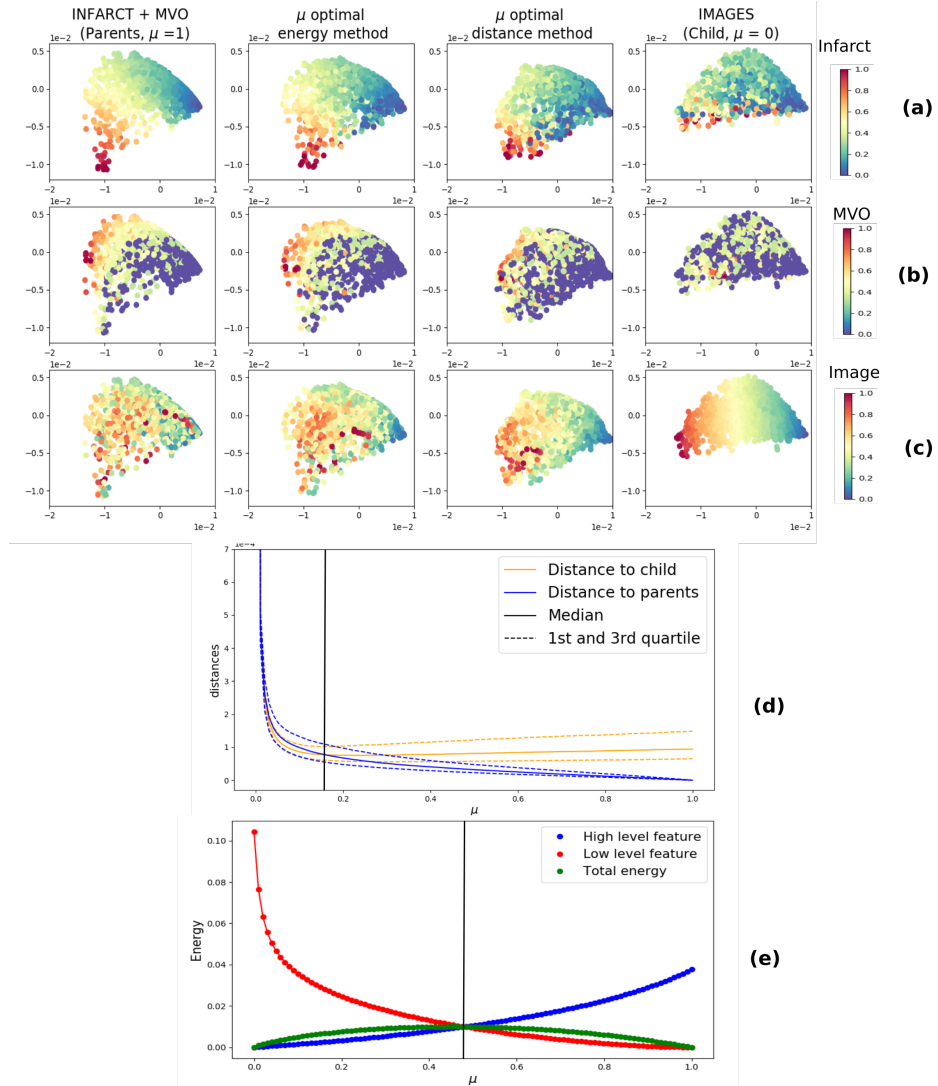


Fig. 2: First three rows (a-b-c): Latent spaces obtained for the parent (left column) and child (right column) data independently, and for the optimal hierarchical embedding (middle columns). The first two dimensions are displayed, and colored according to the infarct (a), MVO size (purple dots correspond to slices without MVO) (b) or the average pixel value in the images (c). Last two rows (d-e): The two optimal μ (vertical black line) were obtained from the crossing of the distance curves (d) or energy curves (e), as explained in Section 2.4.

affected by one specific sample. Conversely, if σ is too large, it can lead to a shrinkage of the latent spaces to zero. We therefore heuristically set σ to the lowest value allowing a meaningful data spread in the latent space (in our case $\sigma = 35$ for the segmentation embedding, and $\sigma = 25$ for the image-based one).

As visible in the two central columns of Fig. 2a-b, the optimal values for μ lead to intermediate latent spaces to which both the parent and child levels contribute. Samples are no longer fully organized according to the segmentation characteristics (the amount of infarct and MVO encoded in the colorscale of Fig. 2a-b). They are neither strongly disorganized as obtained with the child data only, but contain part of the information on image appearance (the average pixel value is encoded in the colorscale of Fig. 2c).

The energy and distance curves in Fig. 2d-e confirm that $\mu = 1$ leads to matching the parent data embedding. In contrast, a jump is observed when approaching $\mu = 0$, as Eq. 3 would amount to $\mathbf{Y}^1 = 0$. This might be explained by the sizes of the embeddings we obtained when approaching $\mu = 0$ (typically for $\mu \leq 0.1$). As the embeddings are really small (due to numerical issues), their rescaling might be less accurate, leading to the jump in the distance curves. On this population, a value of $\mu = 0.47$ leads to balanced energies in Eq. 3, although a smaller value of 0.16 is needed to get a hierarchical embedding equally distant from the child and parent embeddings. In our case, the energy-based solution is closer to the parent embedding structure, whereas the distance-based one is closer to the child latent space.

3.3 Neighborhood consistency

Figure 3 complements these observations by showing representative cases picked from the latent spaces. Row (a) shows a slice with a standard infarct pattern containing a small MVO. The four images and segmentations displayed on the left of the figure stand for the four nearest neighbors in the parent embedding, whereas the ones on the right stand for the four neighbors from the hierarchical embedding (balanced energies solution). We observe that both the neighbors from the parent and hierarchical embedding have segmentations close to the sample slice. However, the images from the hierarchical method are closer (compared with the parent embedding) to the original subject. It means that the embedding is more faithful to the MR image appearance. Row (b) displays the neighbors of a slice with a large MVO that covers most of the infarct. In this case, we display on the left the four nearest neighbors from the image-based embedding, while the images on the right correspond to the neighbors from the hierarchical embedding. We observe here that the image-based neighbors have very different MVO patterns, despite close image appearance. In contrast, the hierarchy-based neighbors are much more consistent with the original pattern. This highlights the interest of guiding the hierarchical embedding by the segmented data, which leads to embeddings more robust to challenging image contents for unusual samples.

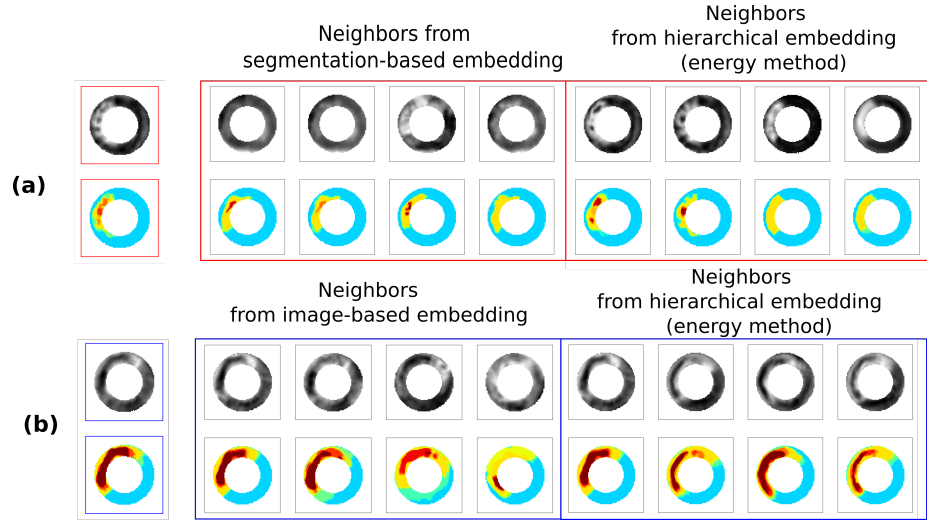


Fig. 3: Illustration of the robustness to challenging image contents. The first column represents the segmentation and image contents for two representative samples: a standard infarct pattern with small MVO (a), and a large MVO covering most of the infarct (b). The remaining columns show their four nearest neighbors for these two cases, either picked from the segmentation- or image-based embeddings, compared to the ones from the hierarchical embedding.

4 Discussion

We have demonstrated the relevance of a hierarchical approach for the analysis of tissue heterogeneities in LGE images. Our hierarchy consisted in guiding the representation of a higher-level (and more challenging) content, by a lower-level (and easier to represent) content corresponding to the segmented images. We also introduced two ways of selecting a relevant weighting parameter μ to balance the contribution of each level in the hierarchy.

Our approach comes from the hierarchical way physicians integrate several data from different sources, for more robustness and confidence in their diagnosis. In our case, the hierarchy allows the integration of prior knowledge corresponding to the lesion location. It can be seen as a way to circumvent the simplicity of a pixelwise distance metric between samples, which may be confounded by specific lesion patterns or image artifacts. We used only simple metrics (pixelwise distances) and a well-known manifold learning framework (diffusion maps), which may reveal limited for complex datasets. Our experiments demonstrated that meaningful embeddings can be achieved even for the challenging cases included in our database.

This hierarchical manifold learning algorithm could be extended to other imaging protocols featuring several levels of data complexity (e.g. typical echocardiography or MRI examinations). The latent space can serve as an intermediate representation to feed a classification algorithm, or as a simplified way to exam-

ine complex data, as in our application. It allows the extraction of the principal diffusion dimensions in such data spaces.

Better image metrics and more powerful algorithms could improve our results in the future, but our purpose here was to demonstrate the interest of such a hierarchy on a sound application. Future work will include the integration of complementary imaging of the lesions such as T1 native and T2* images, early gadolinium enhancement images, and myocardial deformation extracted from image sequences, with the purpose of better understanding ischemia-reperfusion mechanisms in the case of acute myocardial infarction.

Acknowledgements: The authors acknowledge the support from the French ANR (LABEX PRIMES of Univ. Lyon [ANR-11-LABX-0063] and the JCJC project “MIC-MAC” [ANR-19-CE45-0005]), and the Fédération Française de Cardiologie (“MI-MIX” project, Allocation René Foudon).

References

1. Alexandre, J., et al.: Scar extent evaluated by late gadolinium enhancement CMR: a powerful predictor of long term appropriate ICD therapy in patients with coronary artery disease. *J Cardiovasc Magn Reson* **15**, 12 (2013)
2. Bekkers, S., et al.: Microvascular obstruction: underlying pathophysiology and clinical diagnosis. *J Am Coll Cardiol* **55**, 1649–60 (2010)
3. Belle, L., et al.: Comparison of immediate with delayed stenting using the Minimalist Immediate Mechanical Intervention approach in acute ST-segment-elevation myocardial infarction: the MIMI study. *Circ Cardiovasc Interv* **9**, e003388 (2016)
4. Bhatia, K., et al.: Hierarchical manifold learning for regional image analysis. *IEEE Trans Med Imaging* **33**, 444–61 (2014)
5. Bulluck, H., et al.: Cardiovascular magnetic resonance in acute st-segment-elevation myocardial infarction: recent advances, controversies, and future directions. *Circulation* **137**, 1949–64 (2018)
6. Coifman, R., S, L.: Diffusion maps. *Appl Computational Harm Anal* **21**, 5–38 (2006)
7. Criminisi, A., Shotton, J.: Decision forests for computer vision and medical image analysis. Springer Publishing Company (2013)
8. Lin, Y., et al.: Multiple kernel learning for dimensionality reduction. *IEEE Trans Pattern Anal Mach Intell* **33**, 1147–60 (2011)
9. Mom, K., et al.: Population-based personalization of geometric models of myocardial infarction. *Proc. FIMH* **12738**, 3–11 (2021)
10. Shekelle, P.: Clinical practice guidelines: what’s next? *JAMA* **320**, 757–8 (2018)
11. Wang, B., et al.: Unsupervised metric fusion by cross diffusion. *Proc. CVPR* (2012)