Structured Dictionaries for Ischemia Estimation in Cardiac BOLD MRI at Rest

Cristian Rusu[1] and Sotirios A. Tsaftaris[1]

IMT Institute for Advanced Studies Lucca, Italy,
{cristian.rusu,sotirios.tsaftaris}@imtlucca.it,
http://prian.imtlucca.it

Abstract. Cardiac Phase-resolved Blood-Oxygen-Level-Dependent (CP–BOLD) MRI examines changes in myocardial oxygenation in response to ischemia without contrast and stress agents. Since signal intensity changes are subtle, quantitative approaches are necessary to examine variations in myocardial BOLD signals and identify ischemic myocardial territories. Here, using data from animal studies, we extract myocardial time series (BOLD signal as a function of cardiac phase) and explore such variations using a structured dictionary-learning framework, considering shift-invariant learning and spatial priors. We use it: to learn a model of baseline (absence of disease) myocardial time series; and in datasets where disease is assumed, to obtain a spatial map of ischemia presence, identifying myocardial time series from ischemic territories in an unsupervised fashion, by exploiting structural properties, or the lack thereof, in the data. By providing new visualization and quantification approaches, we hope to accelerate the clinical translation of cardiac BOLD MRI for noninvasive ischemia assessment.

1 Introduction

Cardiac phase-resolved Blood-Oxygen-Level-Dependent (CP–BOLD) MRI captures simultaneously BOLD changes (≈10–20%) and wall motion in a single cine acquisition (2D+time) [1] and can assess myocardial ischemia at rest without contrast and stress agents [2], avoiding the contraindication and risks of the latter [3]. Myocardial signal intensities obtained with CP–BOLD vary in a spatio-temporal manner as shown in Fig. 1: when viewed as a 1D time series, BOLD signal intensity remote from the ischemic territory varies across the cardiac cycle in a structural fashion appearing maximal in systole and minimal in diastole, while within the ischemic territory such structure is not present.

These subtle temporal variations make visual assessment challenging (Fig. 1). To aid diagnosis, in [2] ischemic segments (following the classical AHA segmental analysis) were identified if the ratio of average segmental intensity obtained from the images at end-systole (ES) and diastole (identified based on ventricular blood volume) was less than one. However, as [4] suggests, time series obtained from different but healthy territories can appear shifted (in time), thus maximal signal intensity may not always occur in ES, which will affect the performance of [2].

P. Golland et al. (Eds.): MICCAI 2014, Part II, LNCS 8674, pp. 562–569, 2014.

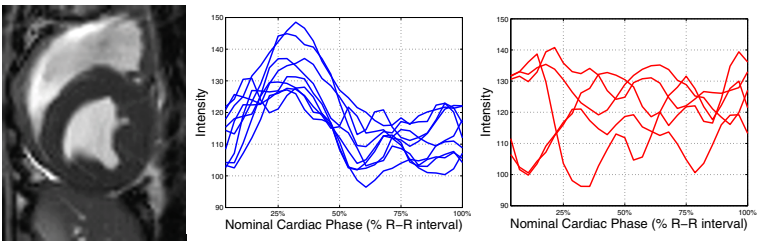

Fig. 1. An image of CP–BOLD MRI at rest in a subject with LAD stenosis (left), and extracted 1D time series from several remote territories (middle) and within the ischemic (LAD) territory (right) as a function of cardiac phase.

In this paper, we propose a new direction in algorithmic design that is able to discern spatio-temporal variations in myocardial time series with robustness to the presence of shifts. It uses all available information in the cine stack (and not only two images as with [2]). By enforcing sparsity, is also robust to noise in the time series attributed to scanning artifacts, physiological variability or inaccuracies in myocardial segmentation and registration.

We rely on structured multi-component dictionary learning framework to model myocardial time series, extracted after segmenting the myocardium and establishing correspondence (via registration) across the cardiac cycle. We decompose each time series as a linear combination of few learned atoms, and group together time series that use the same atoms. To 'guide' the grouping of spatially-neighboring time series we enforce sparsity, spatial priors and explicitly account for the shift structure, by learning atoms using shift-invariant [5] and general [6] dictionaries. Based on the grouping we estimate a probability of belonging to each group. In canine experiments when disease is absent, we find a strong BOLD characteristic behaviour in signal intensities, which varies with cardiac phase and shifts temporally. When ischemia is present, we exploit this behaviour within our dictionary framework to separate myocardial time series by how close they come to this behaviour, and provide a spatial map of ischemia in the myocardium. To aid separation, we also incorporate a discriminatory term, obtaining supervision automatically by looking for said behaviour. The proposed framework jointly optimizes for the unknowns and obtains a final spatial map.

2 Multi-Component Dictionary Learning: MC–DL

2.1 MC–DL for Learning Baseline CP–BOLD Time Series

We consider as input K baseline 2D(+time) CP-BOLD MRI datasets (with no disease present) obtained with a BOLD cine sequence [1]. We want to learn a linear model that describes patterns and statistics in these datasets.

Pre-processing. For each subject k, we trace endo- and epi- cardial boundaries in one image (in systole) and divide the myocardium into $j = 1, ..., J$ non-overlapping consecutive radial segments. We propagate and track these segments

automatically to all images of the cardiac cycle [7]. Thus, for each dataset we obtain J 1D time series reflecting the average BOLD signal intensity of each segment as a function of cardiac phase. We spline interpolate all time series to length n across the study population. To remove any variability due to differences in ES, we estimate for each subject ES as the minimum of ventricular blood volume (from the myocardial delineations) and shift their time series to align them. Finally, we obtain an input matrix $\boldsymbol{Y} \in \mathbb{R}^{n \times K \cdot J}$ containing the columnwise concatenation of all time series in the study population, arranging in the same order of extraction first the time series of the first subject, and so on.

Dictionary learning. We factorize $\boldsymbol{Y} \approx \boldsymbol{D}\boldsymbol{X}$ following the formulation called Multi-Component Dictionary Learning (MC-DL):

$$\begin{aligned} \underset{\boldsymbol{C},\boldsymbol{G},\boldsymbol{X}^{(c)},\boldsymbol{X}^{(g)}}{\text{minimize}} \quad & \left\| \boldsymbol{Y} - \begin{bmatrix} \boldsymbol{C} & \boldsymbol{G} \end{bmatrix} \begin{bmatrix} \boldsymbol{X}^{(c)} \\ \boldsymbol{X}^{(g)} \end{bmatrix} \right\|_F^2 + \alpha \Omega(\boldsymbol{X}) + \beta \Phi(\boldsymbol{X}, \boldsymbol{D}) \\ \text{subject to} \quad & \|\boldsymbol{x}_i^{(c)}\|_0 = 1,\ 1 \le i \le K \cdot J,\ \|\boldsymbol{d}_j\|_2 = 1,\ 1 \le j \le p+g, \end{aligned} \tag{1}$$

where $\boldsymbol{D} = [\boldsymbol{C} \ \ \boldsymbol{G}] \in \mathbb{R}^{n \times m}$ is called *dictionary* (and its normalized columns are called *atoms*) and $\boldsymbol{X} = [\boldsymbol{X}^{(c)};\ \boldsymbol{X}^{(g)}] \in \mathbb{R}^{m \times K \cdot J}$ is the sparse representation matrix with regularization parameters $\alpha, \beta > 0$. Operator Ω imposes sparsity on the columns of $\boldsymbol{X}$, while Φ adds spatial constraints either on the reconstruction $\boldsymbol{D}\boldsymbol{X}$ or on the presentation pattern from $\boldsymbol{X}$.

To accommodate for (shift-invariant) structure in the input time series we construct a dictionary with specific properties. We split the dictionary into a circulant component ($\boldsymbol{C} \in \mathbb{R}^{n \times p}$) and a general component ($\boldsymbol{G} \in \mathbb{R}^{n \times g}$) resulting in the overall dictionary $\boldsymbol{D} = [\boldsymbol{C} \ \ \boldsymbol{G}]$. The split is equivalent in the representations $\boldsymbol{X} \in \mathbb{R}^{(p+g) \times K \cdot J}, \boldsymbol{X} = [\boldsymbol{X}^{(c)};\ \boldsymbol{X}^{(g)}]$. We model input data assuming that it is dominated by a characteristic pattern that shifts across time. Following this expectation, we model each time series as a linear combination between the pattern (or one of its immediate shifts) and a group of general components that can account for other variability (physiological or not). To solve (1) we iterate a double alternative optimization procedure: we solve for variables grouped in pairs of one dictionary component and its representations (($\boldsymbol{C}, \boldsymbol{X}^{(c)}$) and ($\boldsymbol{G}, \boldsymbol{X}^{(g)}$)) and then for each pair we use again an alternative optimization technique popular in dictionary learning with appropriate dictionary updates (circulant or general). The normalization constraint applies to all dictionary atoms.

Sparsity constraints $\Omega(\bullet)$. To enforce sparsity we consider two sparse approximation approaches: matching pursuit and ℓ_1 minimization. We enforce the hard constraint of sparsity 1 for each column of $\boldsymbol{X}^{(c)}$ by considering the highest correlation between each time series with one of the circulant kernel and for the sparsity of $\boldsymbol{X}^{(g)}$ we use either matching pursuit (OMP [8], S–OMP [9]) or the convex penalty $\Omega(\boldsymbol{X}) = \sum_{i=1}^{K \cdot J} \|\boldsymbol{x}_i^{(g)}\|_1$ based on the ℓ_1 norm [10].

Dictionary updates for $\boldsymbol{G}$ and $\boldsymbol{C}$. The general component $\boldsymbol{G}$ is updated using the K–SVD [6] procedure. In the case of $\boldsymbol{C}$ we use a new circulant dictionary learning algorithm (C–DLA) [5]. We remind the reader that circulant matrices are matrices where all columns are circularly shifted versions of the first column $\boldsymbol{c} \in \mathbb{R}^n$, called kernel. In its original form, C–DLA constructs a full circulant

dictionary $\boldsymbol{C} \in \mathbb{R}^{n\times n}$ that reduces the target objective $\|\boldsymbol{Y} - \boldsymbol{C}\boldsymbol{X}^{(c)}\|_F$. Manipulating $\boldsymbol{X}^{(c)}$ we can construct partial circulant dictionaries, e.g., $\boldsymbol{C} \in \mathbb{R}^{n\times p}$ that contains only $p \leq n$ non-consecutive atoms. These constructions are useful since a) we do not expect input datasets to include all possible shifts of a characteristic, but a localized subset, and b) in many cases two consecutive shifts will convey similar information (they are highly correlated) and thus a particular selection pattern (e.g., only atoms with odd indices) may prove advantageous. Here we deal with $\boldsymbol{X}^{(c)} = \left[\boldsymbol{x}_1^{(c)}\ \boldsymbol{0}\ \boldsymbol{x}_3^{(c)} \ldots \boldsymbol{x}_{2p-1}^{(c)}\ \boldsymbol{0} \ldots \boldsymbol{0}\right]^T$, where the nonzero rows are obtained from an OMP run with only the p atoms and they correspond to p nonconsecutive, odd indexed, circulant atoms.

Spatial constraints $\Phi(\bullet)$. The time series in $\boldsymbol{Y}$ are radially arranged and thus provide spatial information. For each subject separately, we explicitly account for this spatial correlation by using one of two known convex penalty functions:

$$\Phi_1(\boldsymbol{X}) = \sum_{i=1}^{J} \|\boldsymbol{x}_i - \boldsymbol{x}_{i+1}\|_1 \text{ or } \Phi_2(\boldsymbol{X}, \boldsymbol{D}) = \sum_{i=1}^{J} \|\boldsymbol{D}\boldsymbol{x}_i - \boldsymbol{D}\boldsymbol{x}_{i+1}\|_2, \qquad (2)$$

where the indices wrap modulo J. For neighbouring time series, Φ_1 enforces similar sparsity patterns (hence the ℓ_1 constraints) while Φ_2 enforces smooth transitions between them by imposing similarity (in the ℓ_2 norm) in their reconstruction. In the alternating optimization procedure that solves the dictionary learning problem these two penalties are considered when updating the representations $\boldsymbol{X}$, thus $\boldsymbol{D}$ is available. Alternatively, we also consider S–OMP [9], which provides sparse approximations with an explicit constraint of equality on the sparsity pattern of predefined groups, which in our case are groups of time series that use the same circulant component among the p.

Subspace identification. After solving the MC–DL problem, we use the representations in $\boldsymbol{X}$ to group and analyse the dataset into subspaces, the unique linear combinations that use the same sparsity pattern (same combinations of atoms). If from the total number of available time series $K \cdot J$ we find N_i having the sparsity pattern of subspace $\mathcal{S}_i$ we attach to this subspace an initial importance of $N_i/(K \cdot J)$. Since the number of possible subspaces is combinatorial we keep the dimensions of the dictionary low (or impose sparsity pattern similarity).

To analyze relationships between the learned subspaces, based on the principal component of each subspace (extracted by the singular value decomposition of the time series in the space) we establish correlations, μ_{ij}, between the subspaces i and j and construct a left stochastic matrix $\boldsymbol{T}$ (i.e., $\sum_i T_{ij} = 1$), with transitions T_{ij} describing the probability of transitioning from subspace i to j:

$$T_{ii} = N_i/(K \cdot J) \text{ and } T_{ij} = |\mu_{ij}|(\sum_{k\neq j} |\mu_{kj}|)^{-1}\,(1 - T_{ii}) \text{ when } i \neq j. \qquad (3)$$

The probability of staying on a subspace (T_{ii}) is defined by its initial importance while the probabilities of transition are the weighted correlations between subspaces. The final probabilities are defined by the stationary vector $\boldsymbol{\pi}$: the left eigenvector of $\boldsymbol{T}$ associated with the eigenvalue of value one.

2.2 MC–DL for Unsupervised Ischemia Estimation in CP–BOLD

We assume as input a single, $K = 1$, CP–BOLD MRI stack of a subject. We extract time series as above, but now $\boldsymbol{Y} \in \mathbb{R}^{n \times J}$ contains only radially arranged time series of one subject. The objective is to find territories of ischemia. We assume that remote to the ischemia time series will behave as in the baseline case. Thus, we separate time series in $\boldsymbol{Y}$ by how much circulant structure they contain, from those that do not, in an unsupervised fashion.

To achieve this we first exploit the result (atoms and initial decomposition) of MC–DL, especially the column-by-column residual energy contained in $\boldsymbol{E} = \boldsymbol{Y} - \boldsymbol{D}\boldsymbol{X}$. Unlike approaches that explicitly add constraints on $\boldsymbol{E}$ [11], we analyse the energy content of the residual with respect to the circulant components learned and the total energy of the dataset. We obtain an initial ischemia estimate by:

$$\mathcal{I}_i = \text{sgn}(\sum_j x_{ji}^{(c)}) \|\boldsymbol{x}_i^{(c)}\|_2 \|\boldsymbol{y}_i\|_2^{-1}, \; i = 1, \ldots, J, \tag{4}$$

As we will see in the next sections, we link high values of this indicator with the behavior of BOLD signal intensities under baseline conditions. The first term in (4) punishes undesirable negative correlations while the other two measure the similarity to the circulant. Based on this initial eschimia estimate $\mathcal{I}_i \in [-1, 1]$, we then deploy label consistent dictionary learning [12] that adds to the overall optimization problem (1) a classification error $\gamma \|\boldsymbol{H} - \boldsymbol{W}\boldsymbol{X}\|_2$, where $\boldsymbol{H} \in \mathbb{R}^{2 \times J}$ contains the class labels obtained by thresholding $\mathcal{I}_i$ and $\boldsymbol{W} \in \mathbb{R}^{2 \times m}$ contains the parameters of the linear classifier. Following the setting of [12] we update all variables, including $\mathcal{I}_i$, in the same alternative iterative fashion. The final estimates are provided by $\mathcal{I}_i$.

3 Results and Discussion

Baseline Model Accuracy: We use $K = 10$ CP–BOLD MRI data obtained from canines with scan parameters as described in [2]. We set $J = 24$ segments, $n = 28$ and follow the pre-processing steps of Sec. 2.1.

With MC–DL we extract, by the circulant component shown in Fig. 2 upper left corner, a *CP–BOLD characteristic pattern* of BOLD variation throughout the cardiac cycle. Fig. 2 further shows the result obtained by the MC–DL with $p = 3$ circulant non-consecutive odd indexed components, $g = 6$ general components for the target sparsity $s = 6$ imposed by the S–OMP algorithm for three groups, corresponding to the three circulant components. The left panel shows the extracted circulant components and dataset partitioned into three groups corresponding to the three allowed shifts and thus the three subspaces (sparsity patterns), with their importance $\boldsymbol{\pi}$ as defined in the end of Sec. 2.1, shown in parenthesis. The right panel shows how important each atom is (energy content on the rows on $\boldsymbol{X}$) and shows clearly the presence of shifts confirming [4]. It also provides a mechanism for selecting p: stop adding circulant atoms when a general one has higher importance. To choose g indicators such as the overall

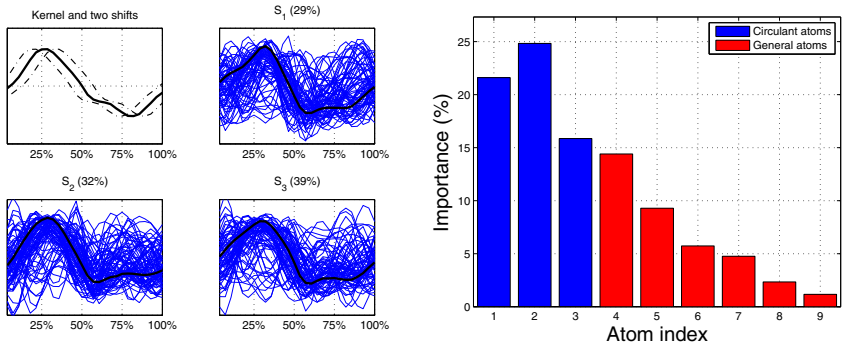

Fig. 2. Model estimated with MC–DL with $p = 3, g = 6$ and target sparsity $s = 6$ using the S–OMP from baseline time series: (left) the circulant kernel, two of its shifts and the three subspaces they construct; (right) relative atom importance.

error ϵ or a low importance can be used. Table 1 shows how different regularizer values affect for MC–DL with ℓ_1 penalty the relative representation error ($\epsilon = \|\boldsymbol{Y} - \boldsymbol{DX}\|_F / \|\boldsymbol{Y}\|_F$, %) and the relative number of nonzero coefficients in the representations (NNZ, %). **Unsupervised Ischemia Estimation:** To test our method, we use an animal model of ischemia. Following [2], we placed a hydraulic occluder in the LAD of a canine to induce controllable stenosis. The animal was scanned twice with CP–BOLD MRI: before stenosis (baseline, absence of disease) to serve as control and at 20 minutes after stenosis (ischemia). Stenosis was held for 3 hours and upon release and reperfusion we obtained a late gadolinium enhancement scan and post-mortem gross histology for visual reference of the infarct. We use the pre-processing steps (Sec. 2.2), with $n = 28$, and $J = 24$ and analyze separately the CP–BOLD data before stenosis and during ischemia. We construct a model based on the MC–DL, ℓ_1 penalty, with dimensions as before ($p = 3, g = 6, \alpha = 1, \beta = 1$) and use the threshold of 0.5 on $\mathcal{I}_i$ to set $\boldsymbol{H}$. The final values of $\mathcal{I}_i$ are color-coded for visualization and estimation of ischemia. Findings are shown in Fig. 3. In (c) a clear infarct in the LAD territory is seen and confirmed by gross histology (f). With the proposed method processing CP–BOLD data under ischemia (b) we find within the LAD territory, significant deviations (color-coded with red hues) from the CP–BOLD characteristic pattern (green hues). These deviations are not present before stenosis (a). Our method because it jointly considers the correlation in

Table 1. Relative representation error and relative number of nonzero coefficients (ϵ / NNZ) for various regularization parameters α, β and penalty Φ_2.

$\alpha \setminus \beta$	0	0.5	1	5	10
0	12% / 100%	12% / 100%	13% / 100%	16% / 100%	22% / 100%
0.5	12% / 93%	13% / 94%	13% / 94%	16% / 95%	22% / 96%
1	13% / 86%	14% / 88%	14% / 88%	18% / 90%	24% / 90%
10	31% / 31%	31% / 31%	32% / 31%	36% / 33%	42% / 33%

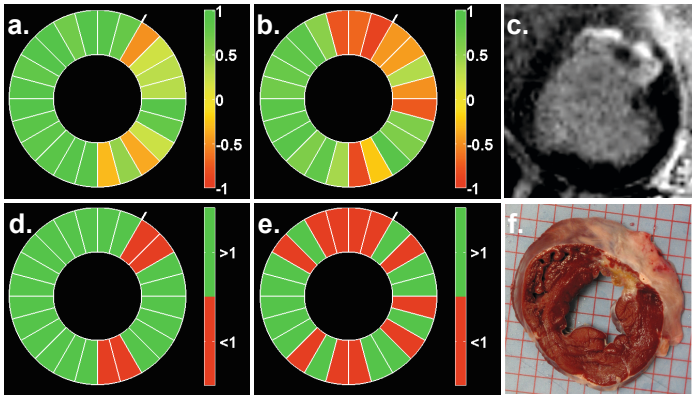

Fig. 3. Comparison of the proposed spatial maps with [2] on a canine with LAD stenosis: green no ischemia present, red ischemia.

the dataset under sparsity and spatial constraints, it obtains accurate and realistic continuous territories of ischemia presence, and reduces false positives when disease is absent. On the contrary, as shown in (e) the approach in [2], due to the hard threshold (> 1: green; < 1 red), finds erroneously ischemia throughout the myocardium and also when disease is absent (d). Relying only on two values at end-systole and end-diastole, it is more susceptible to shifts and noise in the time series (possibly from artifacts or errors in registration).

In the future, to enable true pixel-level analysis, we will incorporate more precise myocardial segmentation [13, 14] and registration [15, 16]. Also, we plan to add a probabilistic interpretation to the spatial maps, as commonly performed in brain fMRI, exploiting further the statistics of the learned spaces in baseline.

This paper uses 2D+time imaging and as such longitudinal motion can influence the findings, which can be seen as limitation. The desirable 3D+time imaging requires prohibitive long breath holds, which can be solved by free-breathing approaches under development. We do not provide fully quantitative validation, but a proof of concept. Simultaneous PET-MR experiments, under development, can provide the means for such validation, with PET as the gold-standard.

4 Conclusion

In this paper we propose a structured multi-component dictionary learning framework to learn and extract characteristic behaviour in myocardial BOLD signal intensity time series as obtained with CP–BOLD MRI. We use circulant and general dictionaries to decompose the data and spatial priors to enforce the model to learn local correlations. We use the same framework to learn how data behave when disease is absent and to identify in an unsupervised fashion ischemic territories in the myocardium in datasets when ischemia is present. We tested our approach on data from canine experiments and found that it outper-

forms previous approaches, opening a future towards a pixel-level visualization and quantification of ischemia on the basis of CP–BOLD MRI at rest.

Acknowledgment. This work was supported in part by the US National Institutes of Health (2R01HL091989-05).

References

1. Zhou, X., Tsaftaris, S.A., Liu, Y., Tang, R., Klein, R., Zuehlsdorff, S., Li, D., Dharmakumar, R.: Artifact-reduced two-dimensional cine steady state free precession for myocardial blood-oxygen-level-dependent imaging. JMRI 31(4), 863–871 (2010)
2. Tsaftaris, S.A., Zhou, X., Tang, R., Li, D., Dharmakumar, R.: Detecting myocardial ischemia at rest with Cardiac Phase-resolved Blood Oxygen Level-Dependent Cardiovascular Magnetic Resonance. Circ. Cardiovasc. Imaging 6(2), 311–319 (2013)
3. U. S. Food and Drug Administration: FDA warns of rare but serious risk of heart attack and death with cardiac nuclear stress test drugs Lexiscan (regadenoson) and Adenoscan (adenosine) (2013)
4. Ootaki, Y., Ootaki, C., Kamohara, K., Akiyama, M., Zahr, F., Kopcak Jr., M.W., Dessoffy, R., Fukamachi, K.: Phasic coronary blood flow patterns in dogs vs. pigs: an acute ischemic heart study. Med. Sci. Monit. 14(10), 193–197 (2008)
5. Rusu, C., Dumitrescu, B., Tsaftaris, S.A.: Explicit shift-invariant dictionary learning. IEEE Sig. Proc. Let. 21(1), 6–9 (2014)
6. Aharon, M., Elad, M., Bruckstein, A.: K–SVD: An algorithm for designing overcomplete dictionaries for sparse representation. IEEE TSP 54(11), 4311–4322 (2006)
7. Tsaftaris, S.A., Andermatt, V., Schlegel, A., Katsaggelos, A.K., Li, D., Dharmakumar, R.: A dynamic programming solution to tracking and elastically matching left ventricular walls in cardiac cine MRI. In: ICIP, pp. 2980–2983 (2008)
8. Pati, Y., Rezaiifar, R., Krishnaprasad, P.: Orthogonal matching pursuit: recursive function approximation with application to wavelet decomposition. In: Asilomar Conf. on Signals, Systems and Comput., vol. 1, pp. 40–44 (1993)
9. Tropp, J.A., Gilbert, A.C., Strauss, M.J.: Simultaneous sparse approximation via greedy pursuit. In: Proc. ICASSP, pp. 721–724 (2005)
10. Tibshirani, R.: Regression shrinkage and selection via the lasso. Series B 58(1), 267–288 (1996)
11. Adler, A., Elad, M., Hel-Or, Y., Rivlin, E.: Sparse coding with anomaly detection. In: Proc. IEEE MLSP (2013)
12. Jiang, Z., Lin, Z., Davis, L.S.: Label consistent K-SVD: Learning a discriminative dictionary for recognition. IEEE PAMI 35(11), 2651–2664 (2013)
13. Feng, C., Li, C., Zhao, D., Davatzikos, C., Litt, H.: Segmentation of the left ventricle using distance regularized two-layer level set approach. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part I. LNCS, vol. 8149, pp. 477–484. Springer, Heidelberg (2013)
14. Bai, W., Shi, W., O'Regan, D.P., Tong, T., Wang, H., Jamil-Copley, S., Peters, N.S., Rueckert, D.: A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: Application to cardiac MR images. IEEE Trans. Med. Imaging 32(7), 1302–1315 (2013)
15. Sundar, H., Litt, H., Shen, D.: Estimating myocardial motion by 4d image warping. Pattern Recogn. 42(11), 2514–2526 (2009)
16. Ou, Y., Ye, D.H., Pohl, K.M., Davatzikos, C.: Validation of DRAMMS among 12 popular methods in cross-subject cardiac MRI registration. In: Dawant, B.M., Christensen, G.E., Fitzpatrick, J.M., Rueckert, D. (eds.) WBIR 2012. LNCS, vol. 7359, pp. 209–219. Springer, Heidelberg (2012)