Synthetic Generation of Myocardial Blood–Oxygen-Level-Dependent MRI Time Series Via Structural Sparse Decomposition Modeling

Cristian Rusu, Rita Morisi, Davide Boschetto, Rohan Dharmakumar, and Sotirios A. Tsaftaris*

Abstract—This paper aims to identify approaches that generate appropriate synthetic data (computer generated) for cardiac phase-resolved blood-oxygen-level-dependent (CP-BOLD) MRI. CP-BOLD MRI is a new contrast agent- and stress-free approach for examining changes in myocardial oxygenation in response to coronary artery disease. However, since signal intensity changes are subtle, rapid visualization is not possible with the naked eye. Quantifying and visualizing the extent of disease relies on myocardial segmentation and registration to isolate the myocardium and establish temporal correspondences and ischemia detection algorithms to identify temporal differences in BOLD signal intensity patterns. If transmurality of the defect is of interest pixel-level analysis is necessary and thus a higher precision in registration is required. Such precision is currently not available affecting the design and performance of the ischemia detection algorithms. In this work, to enable algorithmic developments of ischemia detection irrespective to registration accuracy, we propose an approach that generates synthetic pixel-level myocardial time series. We do this by 1) modeling the temporal changes in BOLD signal intensity based on sparse multi-component dictionary learning, whereby segmentally derived myocardial time series are extracted from canine experimental data to learn the model; and 2) demonstrating the resemblance between real and synthetic time series for validation purposes. We envision that the proposed approach has the capacity to accelerate development of tools for ischemia detection while markedly reducing experimental costs so that cardiac BOLD MRI can be rapidly translated into the clinical arena for the noninvasive assessment of ischemic heart disease.

Index Terms—Blood–oxygen-level-dependent, heart, magnetic resonance imaging (MRI), shift-invariance, sparse decomposition, synthetic generators.

I. INTRODUCTION

S YNTHETIC data (i.e., computer generated data of good resemblance to real data) are used extensively in medical imaging for the development, testing, and validation of image

Manuscript received February 13, 2014; accepted March 12, 2014. Date of publication March 21, 2014; date of current version June 27, 2014. This work was supported by the National Institutes of Health under Grant 2R01HL091989-05. *Asterisk indicates corresponding author.*

C. Rusu is with the IMT Institute for Advanced Studies Lucca, 55100 Lucca, Italy (e-mail: cristian.rusu@imtlucca.it).

R. Morisi and D. Boschetto are with the IMT Institute for Advanced Studies Lucca, 55100 Lucca, Italy (e-mail: rita.morisi@imtlucca.it; davide.boschetto@imtlucca.it).

R. Dharmakumar is with the Biomedical Imaging Research Institute, Cedars-Sinai Medical, Los Angeles, CA 90048 USA (e-mail: rohan.dhar-makumar@cshs.org).

*S. A. Tsaftaris is with the IMT Institute for Advanced Studies Lucca, 55100 Lucca, Italy, and also with the Departments of Radiology, Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208 USA (e-mail: s.tsaftaris@imtlucca.it).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org

Digital Object Identifier 10.1109/TMI.2014.2313000

analysis tasks. In cardiac imaging for example, synthetic data have been used to develop and validate 4-D segmentation [1] and nonrigid registration methods [2], [3]. In functional magnetic resonance imaging (fMRI) of the brain, testing with synthetic data, is common and has helped the broad adoption of certain analysis approaches [4]. The goal of this paper is to identify mechanisms that permit the generation of synthetic data for cardiac phase-resolved blood–oxygen-level-dependent (CP-BOLD) MRI. Using this mechanism we further our understanding and we can generate synthetic data sets to build, test and validate algorithmic developments of ischemia detection.

CP-BOLD MRI is a new contrast agent- and stress-free imaging technique for the assessment of myocardial ischemia at rest [5], capturing simultaneously BOLD changes ($\approx 10\%$ -20%) and wall motion in a single acquisition that can be viewed as a movie (cine BOLD) [6], [7]. Identifying ischemic territories at rest is important, since provocative stress can be contraindicated [8], carrying rare but serious risks [9]. This identification is possible due to differential variations in BOLD signal intensities between systole and diastole, as Fig. 1 illustrates, in a manner which is dependent on the patency of the coronary artery supplying the myocardial territories [5].

With the naked eye such differential variations are not visible thus requiring postprocessing to visualize and quantify the extent of disease. When transmurality of the defect is of interest, pixel-level, rather than segmental-level, analysis can ensure spatial specificity [10]. This analysis entails the following sequence of steps: 1) myocardial segmentation [11]–[16] and subsequent 2) registration [5], [17]–[20] to establish myocardial correspondence at the pixel-level across the cardiac cycle and extract time series of the BOLD variation for each myocardial pixel at each cardiac phase; and 3) an ischemia detection algorithm that is able to discern differences in time series originating from pixels belonging to affected or remote territories, and thus inform the clinician on the part of the myocardium affected by stenosis.

Precise myocardial segmentation and registration still remain a challenge particularly in CP-BOLD imaging due to the temporal variations of the BOLD contrast. The extracted time series can be noisy and affect our understanding of myocardial CP-BOLD effects. The remedy of segmental averaging in turn reduces the number of available time series. Unfortunately, the quality and number of available time series directly affects the design of ischemia detection algorithms. For example, the approach in [5] used only six myocardial segments and two cardiac phases to define a hypothesis-based biomarker. Because of

0278-0062 © 2014 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications standards/publications/rights/index.html for more information.

these dependencies, developments in ischemia detection algorithms are limited and automated pixel-level assessment of ischemia at rest with CP-BOLD imaging is currently lacking.

Synthetic data, when used in lieu of experimental data, can relax such dependencies and allow for the direct development of ischemia detection tools. However, methods to generate such data for CP-BOLD MRI are lacking and current approaches for generating synthetic data in cardiac MRI cannot be readily deployed. For example, image-based methods [3], [21], [22] do not account for cardiac phase-dependent myocardial signal distribution which is necessary for CP-BOLD MRI. Even physicsbased approaches [23]-[26] will also need to account for the cardiac phase-dependent BOLD effect. Currently, two general states (systole and diastole) [5] with two-pool exchange models [27] have been used but multi-state models (one for each cardiac phase of the cine acquisition in CP-BOLD) are not available. Thus, any of these avenues will indeed require the prior investigation of real CP-BOLD time series to understand how the BOLD effect evolves as a function of cardiac phase and in the presence of disease.

In this paper, to further our understanding of BOLD effects, we follow a different approach: we learn statistical models that describe CP-BOLD myocardial time series, on the basis of real experimental data from an animal model of ischemia. To alleviate issues with the accuracy of myocardial registration we use segmental-based time series to learn two models: one for baseline myocardial status and one to describe time series originating from ischemia territories. We then use these models to generate synthetic datasets for which we can alter by design several of their characteristics. We can examine how different acquisition scenarios (acceleration factors and sampling resolution) and the effects of inaccuracies originating from imprecise registration affect algorithm performance, i.e., the correct identification of the disease status of each pixel. Thus, we bypass the steps of segmentation and registration and we can develop and validate ischemia detection algorithms directly in an unbiased fashion. This is important since in the context of BOLD only indirect approaches have been used thus far, relying either on invasive gold standards (such as microsphere-derived segmental perfusion) or on other imaging modalities [5], [10], [28], [29]. By using synthetic data we avoid unnecessary and costly experimentation (involving real subjects) without compromising the validity of the studies and we can accelerate the clinical translation of CP-BOLD.

Specifically, this paper is the first to take a rigorous approach to further our understanding of myocardial CP-BOLD effects using image processing and pattern recognition methods. A new data-driven learning framework builds a statistical model where the CP-BOLD effect is extracted and isolated, relying only on a few input time series. In turn, this allows us to build a synthetic time series generator. Furthermore, we account, using shift-invariant learning, for cases when time series obtained from various myocardial territories under baseline conditions appear slightly shifted (in time) [5], likely because of physiological differences between territories supplied by different arteries [30]. Central to our approach is a sparse multi-component dictionary based method that decomposes time series as linear combinations of template time series (which are also learned

from the data) and their shifts. The method can be viewed as a combination of subspace identification and statistical modeling using Gaussian mixture models (GMMs) for each of the (low-dimensional) subspaces found. Since all learning models exhibit a trade-off between performance and complexity, we introduce a fitness measure that balances the internal dimensions of the proposed model with the representation performance. Since the final goal is to generate synthetic time series with good resemblance to real data we validate this via supervised classification, an approach not previously considered for this task. To show the benefit of having such generators, as proof of concept, we assess the efficacy of a method previously used in fMRI analysis [4] for unsupervised myocardial ischemia detection and compare it with the biomarker introduced in [5]. We evaluate their performance with regards to the number of time series used and their quality (less or more noisy). Unsupervised ischemia detection outperforms the approach in [5], opening the road to significant advances in algorithmic design.

The paper is organized as follows. Section II presents the methodology for learning, modeling, model selection, and validation. Section III presents our experimental findings and Section IV discusses our findings and conclusions.

II. MATERIALS AND METHODS

We propose a model which 1) learns with few input data, 2) offers low representation error when approximating input time series (good representability) without 3) overfitting (i.e., low model complexity) while 4) extracting components and relative parameters that permit the interpretation of the CP-BOLD effect (interpretability) and 5) the controllable generation of time series (controllability).

Given a set of MRI datasets, after necessary preprocessing (Fig. 2, block A) detailed in Section II-A, we obtain an input matrix Y containing a concatenation of time series (examples of which are shown in Fig. 1). Subsequently, in Section II-B, we propose a dictionary learning method to find a sparse approximation of the data based on the dictionary D and the sparse representations X, i.e., each time series (a column of Y) is decomposed as a linear combination of columns of D, weighted by the coefficients in a row of X. This sparse approximation can decompose the input space as a collection of subspaces of much lower dimension (Fig. 2, block B). Based on this approximation, we fit statistical sub-models (Gaussian mixture models) on the distribution of the coefficients of each subspace, i.e., the coefficient blocks of X (Fig. 2, block C). The outcome of this process, called Dictionary Based Modeling Algorithm (DBMA), is a model composed of the dictionary D, the statistical sub-models extracted from X, and its control parameters. Depending on the input data (baseline or ischemic), models are learned for each conditional state.

A learned model (D, M) and a set of parameters (\mathcal{P}) are given to the proposed generator to produce on request synthetic CP-BOLD time series (Fig. 3).

A. Experimental Data and Preprocessing

We use imaging data obtained from a canine experimental model [5] of early ischemia and reperfusion injury briefly

described below. Canines (K = 10) had a surgically implanted hydraulically-activated occluder affecting the left anterior descending (LAD) artery. While anesthetized and mechanically ventilated, the subjects were imaged using a clinical 1.5T MRI system during baseline conditions (i.e., without any occluder activation) and severe stenosis (>90% LAD occlusion). As detailed in [5], CP-BOLD data were acquired mid-ventricle at rest with breath holding twice (pre- and 20 min post-occlusion) using a flow compensated steady state free procession (SSFP) BOLD cine sequence [7] with parameters: field of view, $240 \times 145 \text{ mm}^2$; spatial resolution, $1.2 \times 1.2 \times 8 \text{ mm}^3$; readout bandwidth, 930 Hz/pixel; flip angle, 70°; TR/echo time (TE), 6.2/3.1 ms; temporal resolution 37.2 ms; and 24 time frames. Late gadolinium enhancement (LGE) images were acquired within 20 min post-occlusion (to rule out early infarction) and after 3 h of occlusion and during reperfusion (to identify myocardial regions succumbed to ischemic tissue damage), using the sequence described in [31] with parameters: spatial resolution, $1.3 \times 1.3 \times 8 \text{ mm}^3$; TE/TR, 3.9/8.2 ms; TI, 200 to

We preprocess BOLD stacks to obtain CP-BOLD time series for training and testing the models presented herein (see also Fig. 2, block A). For each subject k out of total K, endo- and epicardial boundaries are manually traced in one image (in systole) and propagated automatically to all images of the cardiac cycle [32]. Subsequently, the myocardium is further segmented automatically into j = 1, ..., J segments/image. For each segment the average intensity is recorded, resulting in average segmental BOLD SSFP signal intensity as a function of cardiac phase, i.e., a time series. All time series are spline interpolated and sampled at n equally spaced time points to create series of the same length across the study population.

220 ms; flip angle, 25°; and readout bandwidth, 140 Hz/pixel.

Each *j*th myocardial segment of the CP-BOLD acquisition at 20 min post-occlusion is identified as "ischemic" if tissue damage is observed on the LGE images acquired after 3 h of occlusion and during reperfusion and as "baseline" otherwise. All segments of the CP-BOLD acquisition pre-occlusion are identified as "baseline." Overall, N time series each of length n and of the same status are concatenated to form the input matrix $Y \in \mathbb{R}^{n \times N}$.

B. MC-DLA: Multi-Component Dictionary Learning

Finding multivariate models (at the extreme of one variable for each of the n cardiac phases) when relying on few training data is difficult due to the known "curse of dimensionality" [33]. Furthermore, as Fig. 1 shows, CP-BOLD time series contain a structure which appears to shift through cardiac phase locations (i.e., the curves appear to be shifted versions of each other). Thus, the learned model should reveal (and retain) such structure to aid interpretation. Various ensemble methods, such as GMMs or random density forests [34], even when combined with dimensionality reduction techniques are agnostic to shiftinvariant structure of the time series and as such do not permit the interpretation of cardiac-phase dependence of the myocardial BOLD effect. On the other hand, due to the sparsity constraint and the robustness to noise [35], sparse decompositions



Fig. 1. Time series extracted from rest CP-BOLD MRI, with a single coronary stenosis, from several remote territories (left) and within the ischemic territory (right) of a single subject.

have been shown to be crucial for the development of robust learning algorithms in high dimensions [36], [37], such as the time series modeling considered herein. Furthermore, sparse decompositions generally offer useful interpretations of the data [38] and based on recent developments can also be handle shifts [39], [40].

Dictionary learning algorithms (DLAs) find dictionary D and sparse representations X such that $Y \approx DX$. Solving this problem involves the ℓ_0 pseudo-norm $\| \bullet \|_0$, the number of nonzero elements. The goal is to minimize, under sparsity constraints, the representation error $\| E \|_F^2 = \sum_{i,j} e_{ij}^2$, i.e., the Frobenius norm of E = Y - DX. However, general DLAs cannot account explicitly for a specific structure [41], as opposed to recent shift-invariant DLA [40]. Combining the advantages of both, here we propose a multi-component dictionary learning method (Fig. 2).

We formulate the multi-component dictionary learning as

$$\begin{array}{l} \underset{\boldsymbol{D}=[\boldsymbol{C}, \boldsymbol{D}_{\boldsymbol{g}}], \boldsymbol{X}}{\text{minimize}} & \|\boldsymbol{Y} - \boldsymbol{C}\boldsymbol{X}^{(C)} - \boldsymbol{D}_{\boldsymbol{g}}\boldsymbol{X}^{(D)}\|_{F} \\ \text{subject to} & \|\boldsymbol{x}_{i}^{(C)}\|_{0} \leq s_{c}, \|\boldsymbol{x}_{i}^{(D)}\|_{0} \leq s_{d} \\ & \|\boldsymbol{d}_{j}\|_{2} = 1, \ 1 \leq j \leq p+m. \end{array}$$

At the core of the proposed formulation is the circulant dictionary $C \in \mathbb{R}^{n \times p}$ with p allowed shifts $(0 \le p \le n)$, where each of them is a down circular-shifted version of the first column denoted as $c \in \mathbb{R}^n$, referred to as *kernel*. Using this kernel we account efficiently for the shift structure in the data. The kernel in baseline conditions, as we will subsequently see, corresponds to the *characteristic CP-BOLD curve* and overall should follow the cardiac phase-dependent variability illustrated in Fig. 1. Its data-driven recovery is a significant innovation of this work.

The general dictionary component $D_g \in \mathbb{R}^{n \times m}$ has m columns d_j called *atoms*. Overall we denote $D = [C, D_g] \in \mathbb{R}^{n \times (p+m)}$ as the multi-component dictionary and $X = [X^{(C)}; X^{(D)}] \in \mathbb{R}^{(p+m) \times N}$, with each column i partitioned in $x_i^{(C)} \in \mathbb{R}^p$ and $x_i^{(D)} \in \mathbb{R}^m$, $1 \le i \le N$. Variables s_c and s_d impose the target sparsity for the circulant and the general dictionary coefficients, respectively. Restrictions on these parameters are naturally set according to the dimensions of each component: $0 \le s_c \le p$ and $0 \le s_d \le m$. The total target sparsity is $s = s_c + s_d$. Dictionaries for which n < mare called overcomplete. Since we want to avoid overfitting our model, in this paper we consider dictionaries with m < n. Algorithm 1 MC-DLA: Multi-Component Dictionary Learning Algorithm. Given the time series Y and dimensions p, m, s_c and s_d , design the dictionary D, and the representations X, such that $||Y - DX||_F$ is reduced.

- Initialization: Compute [C, X^(C)] = C-DLA(Y, p, s_c).
 Iterations: For iter = 1,..., 100
- 1) Compute $[\boldsymbol{D}_{\boldsymbol{g}}, \boldsymbol{X}^{(D)}] = \text{K-SVD}(\boldsymbol{Y} \boldsymbol{C}\boldsymbol{X}^{(C)}, m, s_d).$ 2) Compute $[\boldsymbol{C}, \boldsymbol{X}^{(C)}] = \text{C-DLA}(\boldsymbol{Y} - \boldsymbol{D}_{\boldsymbol{g}}\boldsymbol{X}^{(D)}, p, s_c).$



Fig. 2. Workflow of the proposed learning procedure.

Using the split proposed in (1), the two components are separated and efficient solvers are used to find them. This improves representation performance and interpretability.

To solve (1) we alternately update X, D_g and C. Although this alternate updating is a standard procedure deployed in dictionary learning, where only D_g and $X^{(D)}$ are available [41], here we alternate the three unknowns and use efficient solvers for each [42], [40]. We split the overall target sparsity s among the two components, namely $s_c = 1$ and $s_d = s - 1$.

To recover the general dictionary, the *K-SVD algorithm* [41] uses an alternating optimization procedure to update the unknowns $(D_g, X^{(D)})$: keep D_g fixed and find $X^{(D)}$ using the orthogonal matching pursuit (OMP) algorithm [43]; keep $X^{(D)}$ fixed and update D_g one atom at a time. The power of K-SVD

relies in its atom by atom update rule that simultaneously updates also the corresponding coefficients in the sparse representations. The whole update procedure is based on a rank-1 approximation of an error term. We refer to the K-SVD algorithm as $[D_q, X^{(D)}] = K - SVD(Y, m, s_d)$.

In order to account for the shift-invariant structure we use a recently developed state-of-the-art shift-invariant DLA, the circulant dictionary learning algorithm (C-DLA) [40]. This approach relies on the observation that circulant dictionaries explicitly construct atoms in all their possible shifts. The solution is based on several developments: the eigenvalue factorization of circulant matrices, the internal symmetries of the Fourier transform and the observation that a circulant dictionary C when used together with the OMP algorithm behaves as a shift-invariant system. The final result is obtained by solving a sequence of $\lceil n/2 \rceil$ complex least squares problems of dimension $1 \times N$ involving the Fourier transforms of the rows of Y and $X^{(C)}$. Here, we employ an extension of the C-DLA. Instead of designing a full circulant dictionary C we can create a partial circulant dictionary $C \in \mathbb{R}^{n \times p}$ that allows for only p consecutive shifts explicitly, out of the n total. We refer to C-DLA as $[\boldsymbol{C}, \boldsymbol{X}^{(C)}] = \mathrm{C} - \mathrm{DLA}(\boldsymbol{Y}, p, s_c).$

When put together, K-SVD and C-DLA constitute the proposed MC-DLA, outlined in Algorithm 1. We use a fixed number of iterations, primarily for clarity and because we observed rapid convergence.

With $[D, X] = MC-DLA(Y, p, m, s_c, s_d)$ we refer to MC-DLA. Note that MC-DLA $(Y, 0, m, 0, s_d) \equiv$ K - SVD (Y, m, s_d) and MC-DLA $(Y, p, 0, s_c, 0) \equiv$ C - DLA (Y, p, s_c) . For simplicity, Y is implied and omitted, when allowed.

C. DBMA: The Dictionary Based Modeling Algorithm

Following the previously defined MC-DLA and the general framework presented in Fig. 2, this section describes an efficient sparse generative model of the data. The goal is to create models that allow for a very good representability of the training data, while they retain interpretability and are easy to integrate into a synthetic data generator. To this end, we use the previously defined dictionary and sparse decomposition methods to separate and allow easy access to the feature of interest (i.e., the shift-invariant component), while constructing a natural method to randomly sample from the union of subspaces created by the model (Fig. 2, block B). To achieve this, we statistically model the coefficients used for each subspace and create sub-models that capture a particular behavior in the data (Fig. 2, block C).

The atoms and the subspaces created by their combination, maximum $\binom{m}{s}$ with $m \ge s$, offer deep insight into the geometric properties of the data, while they provide robustness to noise, outliers and overfitting due to the sparsity prior. The method is especially effective in dealing with high dimensional data by projecting results onto lower dimensional subspaces that are easier to manage statistically and algorithmically.

Each sub-model is fully determined by its *s* atoms and an associated Gaussian mixture model (GMM). GMMs have become particularly useful in the modeling of data densities and

Algorithm	2	DBI	MA:	Dictio	nary	Based	Mode	ling	Al-
gorithm.	Give	en t	ime	series	$\boldsymbol{Y},$	create	model	\mathcal{M}	=
$\{\mathcal{M}_1,\ldots,$	\mathcal{M}_L	} that	at acc	urately	desci	ribes the	input.		

- 1) Learn a Gaussian $\mathcal{N}(\mu_m, \sigma_m)$ for time series' mean.
- 2) Normalize time series and remove their mean in Y.
- 3) Invoke [D, X] = MC-DLA (Y, p, m, 1, s 1).
- Find all L subspaces composed of unique combinations of atoms from D using X.
- 5) For each subspace $\ell = 1, \ldots, L$:
 - a) Find time series $Y_{\ell} \in \mathbb{R}^{n \times N_{\ell}}$ used by D_{ℓ} .
 - b) Separate coefficients $X_{\ell} \in \mathbb{R}^{s \times N_{\ell}}$ used by D_{ℓ} .
 - c) Using X_{ℓ} , construct the GMM $\mathcal{G}_{\ell} = \sum_{j=1}^{J_{\ell}} \pi_{j\ell} \mathcal{N}(\mu_{j\ell}, \Sigma_{j\ell})$, where J_{ℓ} is the number of mixtures in the GMM (chosen during training) while $\mu_{j\ell} \in \mathbb{R}^s$ and $\Sigma_{j\ell} \in \mathbb{R}^{s \times s}$ represent the mean vector and covariance matrix of the Gaussian, all created for model ℓ .
 - d) Estimate sub-model \mathcal{M}_{ℓ} importance $P_{\ell} = N_{\ell}/N$.
- 6) Construct the error matrix E and model its contribution as additive white Gaussian noise $\mathcal{N}(\mu_{\epsilon}, \sigma_{\epsilon})$.

have produced satisfactory results in a number of applications [44]. However, in large dimensions they are met with significant difficulty since the limited number of samples results in nonsingular covariance matrices, and thus some form of regularization is required [45]. This motivates our solution: separate the input space into a union of low-dimensional subspaces and statistically model each one.

Algorithm 2 shows the steps of the proposed modeling approach, which we call dictionary based modeling algorithm (DBMA), with a visual description provided in Fig. 2. At the core of DBMA lies the MC-DLA. Based on the inputs p, m and s we create the dictionary D and its sparse representation coefficients X. During the modeling steps the coefficient matrix Xis used exclusively to group the input time series into clusters that lend themselves to easier statistical modeling. Based on X, we search of all subspaces (all unique combinations of s atoms actually used) obtained after the dictionary learning. Albeit the total number of subspaces can be very high, maximum $\binom{m}{s-1}p$ assuming $m \ge s - 1$, we expect to have only a fraction of this number due to highly correlated and structured real world data. To create an effective learning mechanism we use the coefficients as training data to fit GMMs.

For each subspace $D_{\ell} \in \mathbb{R}^{n \times s}$ we use the reduced (without zeros) coefficient matrix $X_{\ell} \in \mathbb{R}^{s \times N_{\ell}}$, where N_{ℓ} denotes the number of time series attributed to sub-model ℓ , to create efficient and stable statistical models on the low-dimensional space of size *s*—note that $s \ll n \ll N$. Since for each D_{ℓ} we know the corresponding atoms from the dictionary, after sampling the GMM we can generate synthetic output data-thus, each sub-model $\mathcal{M}_{\ell} = (D_{\ell}, \mathcal{G}_{\ell})$ contains the subspace information and the statistical GMM associated with it. To estimate the importance of each sub-model P_{ℓ} , we compute the percentage of time series that belong to the subspace. P_{ℓ} can be seen as a probability measure since $\sum_{\ell=1}^{L} P_{\ell} = 1$. Finally, we account for the residual error E as additive Gaussian noise $\mathcal{N}(\mu_{\epsilon}, \sigma_{\epsilon})$.

D. Model Selection

We mentioned previously that the parameters of DBMA must be identified to balance crucial and competing properties: controllability and interpretability against model complexity and representation capabilities. More degrees of freedom (i.e., more atoms and larger target sparsity) lead to lower representation error but they increase model complexity, number of subspaces and statistical sub-models and do risk overfitting. Thus, we should avoid overly complex solutions where each sub-model may reflect idiosyncrasies of the data and not valuable key characteristics. Since we are dealing with different types of components (the *m* generic and the *p* circulant atoms), in order to provide a fair comparison we consider as the total degrees of freedom the maximum number of subspaces the model can reach: $\binom{m}{s-1}p$ in this particular case of MC-DLA.

We propose a measure that can evaluate dimension selection, deviating from traditional performance bounds in sparse decomposition and overcomplete dictionaries. When evaluating the performance of overcomplete dictionaries, in general, two obvious indicators are widely used: the representation performance [41] and the mutual coherence of the dictionary $\mu(D)$ that directly affects the successful recovery rate of sparse approximation algorithms [46]. Both approaches focus on the computation of incoherent dictionary atoms during the training phase [42], [47], while in our case the kernel and its shifted versions may be correlated. Other statistical model selection methods, such as the Akaike or Bayesian information criteria or Minimum description length [48], which aim to balance model likelihood and dimensions, do not fit the approach proposed in this paper since many subspaces are created, each with varying dimension and importance.

Given a model \mathcal{M} , based on MC-DLA(p, m, 1, s - 1) with L sub-models, we introduce the model fitness measure

$$F(\mathcal{M}) = \frac{\epsilon}{L} \binom{m}{s-1} p \tag{2}$$

that not only balances the relative representation capability

$$\epsilon = \|\boldsymbol{Y} - \boldsymbol{D}\boldsymbol{X}\|_F / \|\boldsymbol{Y}\|_F \tag{3}$$

but takes into account also the full available search space (the combinatorial term). The number of sub-models L, relative to the total number of possible sub-models, affects directly the computational complexity of the model.

We are interested in finding the model corresponding to a minimum value of the fitness—low relative representation error and model dimension. Notice that $F(\mathcal{M})$ is dimensionless but positive, with a zero minimum value achieved when the representation error is exactly zero. When m = N a trivial solution exists $(D_g = Y, X^{(D)} = I_N)$ and perfect representation is achieved with target sparsity one. Of course, here we focus on solutions where $m \ll N$.

We will show that given certain conditions, there exist local minima of (2) with respect to the dictionary dimension m. These minima reflect the ideal balance between representation error and complexity. The choice of p will be discussed in



Fig. 3. Proposed generator structure based on DBMA.

Section III-B. Also, for simplicity, but without sacrificing generality, we assume the target sparsity *s* fixed.

Let us consider for a given dataset, fixed sparsity s and shifts p, three models $\mathcal{M}_{(i)}$, i = -1, 0, 1, where each utilizes m + i general dictionary atoms with representation errors $\epsilon_{(i)}$ and $L_{(i)}$ sub-models. Assuming $\mathcal{M}_{(0)}$ is a local minimum

$$F(\mathcal{M}_{(-1)}) \ge F(\mathcal{M}_{(0)}) \le F(\mathcal{M}_{(1)}) \tag{4}$$

using (2) and expanding the factorial terms, we reach

$$\frac{\epsilon_{(0)}}{L_{(0)}} \le \min\left\{\frac{\epsilon_{(1)}}{L_{(1)}}\frac{m+1}{m-s+2}, \frac{\epsilon_{(-1)}}{L_{(-1)}}\frac{m-s+1}{m}\right\}$$
(5)

which shows how representation error and complexity relate for consecutive m. Recall that the problem dimensions obey $1 \le s \le m \ll N$ while $\epsilon_{(i)} \in [0, 1]$ and $L_{(i)} \in \{1, \ldots, N\}$.

Inequality (5) conveys information on the relative importance of representation error and complexity. For example, assuming for simplicity that $L_{(0)} \approx L_{(1)}$ we find that $\epsilon_{(0)}(m-s+2)/(m+1) \leq \epsilon_{(1)}$ must be true in order for $F(\mathcal{M}_{(0)}) \leq F(\mathcal{M}_{(1)})$ to hold. As *m* increases, this bounds the representation error improvement that has to occur for the new model with more atoms to be considered better according to (2).

Generally, we are interested in conditions for which the existence of minima, even local, is possible as $m \ll N$. Result (5) is useful but involves knowing the values of the error and complexity. Although both clearly depend on m, good, generally proven formulations or estimates are not available. Motivated by findings of others [46] and by results in Section III, we study the dynamics of the proposed fitness measure assuming an exponential evolution of the representation error

$$\epsilon_{(i)} = ae^{b(m+i)} \tag{6}$$

with $a, b \in \mathbb{R}, b < 0$ and a > 0. Using these assumptions, writing the complexities $L_{(1)} = \gamma_{(1)}L_{(0)}, L_{(-1)} = \gamma_{(-1)}L_{(0)}$ and including (5), minima points exist if

$$\gamma_{(1)} \le e^b \frac{m+1}{m-s+2}, \ \gamma_{(-1)} \le e^{-b} \frac{m-s+1}{m}.$$
 (7)

In the general case, we can assume that the dynamics of the complexity follows $\gamma_{(1)} \geq 1$ and $\gamma_{(-1)} \leq 1$, i.e., the number of used subspaces increases with the increase in the total search space. These bounds on $\gamma_{(1)}, \gamma_{(-1)}$ are met when $m \leq \lfloor (s - 1)(1 - e^b)^{-1} - 1 \rfloor$. The assumptions and bounds provided in this section fall into the regions of interest of $m \ll N$ and are not particularly restrictive. In Section III we support the choice made in (6) and show experimental results that validate $F(\mathcal{M})$.

E. Sampling From the Sparse Generative Model

As Fig. 3 illustrates, given a learned model \mathcal{M} , the dictionary D and some parameters (denoted as \mathcal{P}), generating \hat{N} synthetic time series can be done efficiently on request since sampling from GMM distributions is well known. First, a sub-model $(D_{\ell}, \mathcal{G}_{\ell})$ is selected and then synthetic time series are generated by drawing projection coefficients following their GMM distribution and using them as weights to linearly combine the atoms of each individual subspace. Concretely, to generate synthetic dataset $\hat{Y} \in \mathbb{R}^{n \times \hat{N}}$ the process is

$$\hat{\boldsymbol{y}}_i = \boldsymbol{D}_{\ell} \boldsymbol{z}_{i\ell} + \boldsymbol{n}_i + m_i \boldsymbol{1}, \ i = 1, \dots, \hat{N}$$
(8)

where $\boldsymbol{z}_{i\ell} \in \mathbb{R}^{s \times 1}, \boldsymbol{z}_{i\ell} \sim \mathcal{G}_{\ell}, \ell$ is chosen randomly with probability $P_{\ell}, \boldsymbol{n}_i \in \mathbb{R}^n, \boldsymbol{n}_i \sim \mathcal{N}(\mu_{\epsilon}, \sigma_{\epsilon}), \mathbf{1} \in \mathbb{R}^n$ is the all-ones vector and $m_i \sim \mathcal{N}(\mu_m, \sigma_m)$. The set \mathcal{P} includes the following.

- 1) \hat{N} —the number of output series to generate.
- 2) S—the sampling method. To generate N time series use: ALL—all of the available sub-models each selected with probability P_{ℓ} or SINGLE—a single sub-model.
- Systolic to diastolic ratio (SDR)—the maximum BOLD contrast, as defined in [5]: ratio between the maximal (endsystole) and last (diastole) values.
- N(μ_ε, σ_ε)—the Gaussian model of the residual. Obtain a vector n_i randomly sampled from the Gaussian.
- 5) $\mathcal{N}(\mu_m, \sigma_m)$ —the Gaussian model of the mean. Obtain a random sample m_i and add it to each time series.

Varying these parameters can reflect many effects such as: different flip angles and magnet strengths (e.g., 1.5T or 3T) [6], accuracy of registration, noise level, resolution and others. While some parameters are modeled in a simple fashion (e.g., as additive Gaussian processes) others stem from the power of the proposed decomposition method to separate effects of interest (BOLD phase dependent variability and shifts) from noise, adding the desired interpretability and controllability. This further permits the definition of a sampling methodology S and the extraction of CP-BOLD related information (e.g., SDR and the characteristic CP-BOLD curve itself).

F. Validation of Synthetic CP-BOLD Myocardial Time Series

For the important task of validation here we indirectly validate the resemblance between synthetic and real data following an application driven approach: we show that synthetic and real data behave the same when used in a supervised learning context. We learn support vector machines (SVM) [49] on the real data (binary classification to baseline or ischemia) and test them on several instances of synthetic data and vice versa.

We utilize the dataset of real K subjects with 2 KJ available time series, for which ground truth is available on the basis of LGE imaging (Section II-A). Keeping the baseline to ischemic time series overall ratio fixed, we use the synthetic generators to create γ sets of K "synthetic subjects," for which ground truth is known by design. We choose an SVM with a radial basis function kernel with scaling factor $\sigma = 1$ and a soft margin $\nu = 1$. We train it using each time series y_i directly as an input vector without any feature extraction.

We consider the following training/testing combinations.

Cross Validation Accuracy on Real Data: We split the real subjects using a leave one subject out strategy, training with K-1 and testing on the remaining one. We measure accuracy as the number of correctly identified time series. The average (and standard deviation) of this cross validation experiment serves as benchmark performance on real data.

Cross Validation Accuracy on Synthetic Data: Similar to the above, for each of the γ sets we train the classifier on the K-1 synthetic subjects, test on the remaining one.

Classification Accuracy of Synthetic With Classifier Trained on Real Data: The SVM is trained with all real K subjects and is tested on each of the γ synthetic sets.

Classification Accuracy of Real Data With Classifier Trained on Synthetic Data: We train the SVM picking one of the γ sets of K synthetic subjects. We test accuracy on the real data.

Accuracy is always the average over all γ sets. An accuracy similar in value across the four combinations, tested with two sample t-tests, would indicate that the classifier performs the same irrespective of training data origin.

III. RESULTS

We utilize imaging data from the same K = 10 canine subjects in controlled rest experiments. Overall, 20 CP-BOLD MRI stacks (two per each subject) are available: 10 without and 10 with critical LAD stenosis. We use J = 24 segments, and interpolate to length n = 28. In total N = 480 segmental time series are available. Based on LGE imaging 3 h post-occlusion and during reperfusion, $N_i = 80$ are considered ischemic (i.e., the segments were affected by the stenosis) and $N_b = 400$ are baseline (i.e., we group together time series from baseline and from remote territories in rest-stenosis studies since they were found to behave the same).

We show in sequence, results on representation accuracy, model selection, DBMA and its validation. Then, we use DMBA in a proof of concept experiment (Section III-D). We conclude showing statistically the existence of shifts (Section III-E).

A. Model Accuracy and Representation Capabilities

To evaluate the representation capabilities of the proposed method we use the metric defined in (3). As reference we include the principal component analysis (PCA).

Referring to Table I, with respect to PCA, the simplest model (the circulant only) performs quite well, indicating a strong circulant structure in the data. By comparison, the general dictionary with unitary sparsity and five atoms offers worse performance. As the number of atoms increases, the general dictionary component outperforms the circulant component, indicating that there is a circulant structure but that it is localized. In fact, observing for MC-DLA(28, 0, 1, 0) which circulant atom (i.e., which shift of the kernel) is mostly used in reconstructing each time series, we can see the importance (energy content) of each shift, shown in Fig. 4. The energy content for each shift is computed as the sum of the squares of the representation coefficients where that shift has the largest correlation (in an OMP fashion, with target sparsity 1). The values are scaled to the total

 TABLE I

 Relative Representation Error ϵ (%) of Various Settings

 OF MC-DLA $(p, m, s_c, s_d), s = s_c + s_d$

Target sparsity (s)/Number of components							
MC-DLA	1	2	3	4	5	6	
	1	Baseline	data				
Circulant only							
(4,0,s,0)	62.9	57.2	44.1	38.8	-	-	
(28, 0, s, 0)	50.6	38.3	36.2	25.5	20.1	17.0	
General only							
(0,5,0,s)	53.4	40.9	34.3	30.6	29.9	-	
$\left(0,10,0,s ight)$	47.3	35.0	27.7	22.3	18.6	16.6	
(0, 15, 0, s)	46.1	30.5	24.3	19.8	17.5	15.3	
Multi-component							
(4, 6, 1, s - 1)	-	36.7	29.0	25.1	22.3	21.5	
(4, 7, 1, s - 1)	_	35.7	27.9	22.8	20.1	18.5	
(4, 10, 1, s - 1)	_	33.7	25.8	20.8	18.1	14.9	
(28, 15, 1, s - 1)	—	28.7	21.3	17.7	14.7	12.5	
PCA	69.9	56.1	43.9	36.2	29.9	25.0	
Ischemic data							
General only							
(0,6,0,s)	50.9	38.0	29.6	25.0	24.1	23.8	
(0,7,0,s)	48.6	35.2	26.8	21.8	19.5	18.7	
(0,8,0,s)	46.7	29.9	24.5	20.1	17.6	16.0	
PCA	78.6	57.5	43.6	35.5	29.1	23.8	
		1	1				



Fig. 4. Global importance of each kernel shift found via the circulant only MC-DLA(28, 0, 1, 0) using all baseline data.

energy content, the square Frobenius norm of the coefficients matrix. The grouped peaks and their wide spread allude to the presence of localized shift structure attributed to variability in ES among subjects.

While for baseline conditions we identify a strong underlying structure, we do not expect or impose such underlying structure in ischemic time series. Thus, we use only the general dictionary component of MC-DLA, to identify possible submodels. Table I depicts representation errors and compares various MC-DLA settings with the PCA. Naturally, the representation error decreases as model complexity increases (model dimensions, p or m, and target sparsity s).

Our findings and the atoms shown in Fig. 5, lend to the following interpretation of the data: under baseline conditions time series contain a strong structure, the *characteristic CP-BOLD curve*, represented here by the circulant kernel and its shifts, plus a collection of additional components (attributed to subject variability, or to registration-related inaccuracies) modeled by the general dictionary atoms, which when linearly combined can alter this curve. The kernel, shown in Fig. 5, relates the intensity of the CP-BOLD signal as a function of cardiac phase: at $\approx 33\%$ of the nominal cardiac cycle the highest value is present. In healthy conditions, this is a physiologically expected location



Fig. 5. Dictionary components (atoms) estimated with MC-DLA(4, 6, 1, 3) from baseline time series: (left) the circulant kernel and its three shifts; (middle) the general dictionary atoms; (right) their relative importance.



Fig. 6. Model fitness for various values of m and p(s = 4).

for end-systole [7], [5] and is in agreement with theoretical findings in [5] that in end-systole BOLD effect is higher.

Throughout this paper we pay attention to extracting the kernel and its shifts, since both are crucial to the interpretability and controllability of the proposed model. We can also use them to recover ES location variability between subjects. In our dataset using the myocardial segmentations (Section II-A), on the basis of the standard identification of ES as the minimum blood volume in the ventricle, we find the location of ES at $37 \pm 6\%$. Observing which kernel is used the most for each subject, combined with the peak's location in the kernel, we estimate an ES average position at $37 \pm 11\%$. A paired t-test comparing the two showed no difference (p - value = 1).

B. Model Selection

Representation performance, albeit crucial, is not the only indicator we are interested in. For example, the model MC-DLA(4, 10, 1, 5) achieves a representation error of 14.9% while the general only model MC-DLA(0, 15, 0, 6) achieves a higher error of 15.3%. Recall that for the multi-component method one component from the total s is restricted to one of the four circulant atoms—compared to the general model that can access 15 independent atoms in any combination.

Since we are not only interested in good representations, we use the proposed fitness measure (2) to search for a more suitable model. For a fixed s = 4, increasing values of m and several values of p, we show in Fig. 6 the fitness of some models. Clearly, a local minimum point is present in m = 6 for p = 4. As we hypothesized, representation error decays exponentially with respect to m as Fig. 7 shows. Following the error model in (6), we estimate $\epsilon(m) \approx 0.21e^{-0.009\text{m}}$.

While in the methods section we focused on the selection of m, with respect to p several observations can be made. Fig. 6 for example shows how with varying p, the fitness value differs,



Fig. 7. Representation error evolution achieved by the general dictionary component, with the increase in m shown on a logarithmic scale, after the removal of the circulant contribution.

and we could repeat the analysis of Section II-D with respect to s. Interestingly, we could identify a suitable value of p via the importance plot in Fig. 5. As p increases the importance for the newly added shift decreases. After a certain cut-off value the importance of the atoms in the general dictionary will outweigh the benefit of further considering shifts of the kernel. This heuristic provides very good results. For example, the p = 4 allowed shifts account for approximately 60% of the energy, showing the strong structure existing in the data. The first noncirculant atom is less important than any circulant one. These experimental verifications motivate the selection of MC-DLA(4, 6, 1, 3) as the model of the baseline time series.

Having demonstrated the representation capabilities of the proposed formulation we are now interested in its modeling and generating capacity. We consider two models of DBMA: $\mathcal{M}_{\rm base}$ for baseline time series (using the composite learning method) found with MC-DLA(4, 6, 1, 3) and $\mathcal{M}_{\rm isch}$ for ischemic time series, which due to the lack of structure uses only general components found via MC-DLA(0, 7, 0, 6).

Fig. 8 shows sub-models of both learned models. For the six baseline sub-models, it is evident that each sub-model represents a different shift of the kernel and alternations to its shape. The percentage of utilization of each sub-model (P_{ℓ}) is also shown. As the utilization decreases the deviation from the characteristic CP-BOLD curve increases. The ischemic sub-models lack a consistent characteristic behavior.

C. Validation and Comparison With a PCA-Derived Model

Since DBMA is designed specifically for BOLD time series analysis and no reference method exists, direct comparisons with other approaches are difficult. General statistical modeling approaches, while they can still provide acceptable representation error (see the performance of PCA in Table I) they do



Fig. 8. Representative sub-models found via MC-DLA(4, 6, 1, 3) for baseline (in blue) and MC-DLA(0, 7, 0, 6) for ischemia (in red). We show real time series belonging to each sub-model and their principal component found via the SVD (thicker line), with the utilization percent of each.

not necessarily provide insight into the structure of the data and extract the key expected feature: the characteristic CP-BOLD curve (see Fig. 5). It is expected that any model of high complexity should be capable to represent the input data with low representation error. Thus, one must take into account not only error but also parsimony (model size), interpretability and controllability in terms of the CP-BOLD effect.

We compare with a PCA-derived decomposition, which we refer to as PCA-GMM, which uses a single model not several sub-models as with DBMA. With six principal components obtained from Y, using the so-called reduced singular value decomposition $Y = U_6 \Sigma_6 V_6^T$, $\Sigma_6 \in \mathbb{R}^{6 \times 6}$, we learn a mixture of 10 Gaussian components for the coefficients $\Sigma_6 V_6^T$ in the basis U_6 . Since the EM algorithm deployed to fit the mixture encounters convergence problems, the results presented are obtained after several runs with different initializations.

The baseline DBMA generator uses the options: $N_b = 18, S = \text{SINGLE}$ while SDR and the Gaussian models of the residual and the mean are estimated by DBMA. The ischemic generator uses the options: $\hat{N}_i = 6, S = \text{ALL}$, other parameters are estimated by DMBA. We create $\gamma = 100$ synthetic sets reflecting K = 10 synthetic subjects, with $\hat{N} = 24$ total time series. Thus, the created synthetic subjects have similar composition to our real subjects. Fig. 9 shows, for visual comparison, baseline time series used for training (\mathbf{Y}) and synthetic $(\hat{\mathbf{Y}})$, equally numbered.

As presented in Section II-F, we use a supervised approach for validating time series generated by DBMA with four combinations of training and testing sets (results in Table II). In all



Fig. 9. Examples of real (left) and synthetic (right) time series.

TABLE II Accuracy (Mean±Standard Deviation) for the Four Setups of SVM-Based Validation of the Synthetic Generators

	Testing set			
Training set	Real	Synthetic		
Real	$83\pm6\%$	$84\pm0.7\%$		
Synthetic	$85\pm0.5\%$	$83\pm3\%$		

cases an accuracy greater than 80% is observed. (Although this accuracy could increase with better choice of features, this is beyond the scope of this paper.) The goal is to provide strong evidence that the synthetic data closely resemble the real ones. We find no statistical significance when comparing average accuracy independently on the origin of the training data and testing data (t-tests, p-values >0.5). This indicates that the classifiers have comparable behavior when trained on the real and tested on the synthetic dataset and vice versa. Thus, we show indirectly the equivalence between the real and synthetic data (i.e., they lie on the same manifold). The synthetic data generated by the PCA-GMM perform similarly to the results presented in Table II, albeit the modeling approaches are different. This behavior is due to the significant variability between baseline and ischemic time timeseries. However, when trained and tested on synthetic data, PCA-GMM achieves an undesirable higher accuracy $(85\pm2\%)$ indicating that PCA-GMM produces synthetic sets will less variability than the real input data. Note that we are not in the pursuit of high accuracy, but in replicating the results obtained with the real dataset.

Next we evaluate the capability of the synthetic data in retaining the shift structure (and its importance function) of the input data. Given the input dataset and its importance function (see Fig. 4), using both methods (MC-DLA and PCA-GMM) we learn and generate several synthetic datasets containing N = Ntime series. For each synthetic dataset we learn a kernel, extract its importance function, as described in Section III-A, and correlate it with the one of the input dataset (i.e., Fig. 4). Over 200 realizations, MC-DLA obtains an average correlation of 92%, and PCA-GMM of 88%. The difference is statistically significant (p-value $< 10^{-6}$; t-test). Synthetic experiments, not shown for brevity, involving larger datasets of greater shift variability (in real data this would occur due to larger variance in ES position, possibly due to arrhythmia) and noise show that this performance gap increases in favor of MC-DLA since it accounts explicitly for shift structure.

TABLE III Comparing Accuracy (%) of ICA and the S/D Biomarker [5] for Different Synthetic Compositions

	$\hat{N} = 300$		$\hat{N} = 150$		$\hat{N} = 24$	
\hat{N}_i/\hat{N}	ICA	S/D	ICA	S/D	ICA	S/D
2/5	93±5*	73±6	91±5*	73±8	85±7*	68±12
1/3	$88 \pm 5^{*}$	75 ± 6	$87 \pm 6^{*}$	74±7	$82 \pm 7^{*}$	72 ± 12
1/4	$80\pm5*$	78±7	79 ± 6	76±7	77 ± 8	75 ± 12
* denotes statistically significant difference (p-value< 0.05) between ICA						
and S/D for the same choice of \hat{N} and \hat{N}_i .						

D. Synthetic Data to Evaluate Unsupervised Ischemia Detection

As a proof of concept, we use synthetic time series obtained from the two generators as inputs, in lieu of real data, to compare the performance of independent component analysis (ICA), an unsupervised pattern recognition method, with the S/D approach in [5] for ischemia detection. ICA is used extensively in fMRI for the detection of activation areas in a subject's brain [4]: it identifies activated areas (i.e., variations in time series) by decomposing the observed fMRI time series in a linear combination of different spatial components. Drawing a parallel between fMRI and CP-BOLD, we explore ICA as an unsupervised method for separating ischemic from baseline CP-BOLD time series extracted from a "synthetic subject's" data. Using synthetic data we know by design the ground truth, i.e., the originating class of each time series; thus, exact validation is possible, making comparisons between ICA and S/D fully quantitative. We show accuracy with respect to varying the number of time series (reflecting variations in spatial resolution), the percentage that are ischemic (reflecting variations in coronary patency) and finally the amount of artificial noise introduced (reflecting errors in registration or different acquisition scenarios, e.g., acceleration factors).

We assume an input observation matrix originating from a single "synthetic" subject. We use the same setup (same models) as before, and sample the baseline and ischemic generators, varying the composition with respect to the amount of baseline (\hat{N}_b) and ischemic (\hat{N}_i) time series, choosing ratios (\hat{N}_i/\hat{N}) , where $\hat{N} = \hat{N}_i + \hat{N}_b$ reflecting realistic situations of single vessel disease. We follow [4] to obtain a single independent component (of zero mean), which carries information on the class of each time series. Since we are not interested in inference, we use a fixed threshold of 0 on the component's values to label each time series and compare it to the ground truth.

Table III summarizes our findings after generating 100 "synthetic" subjects for each composition. To demonstrate the benefit of utilizing all cardiac phases, in Table III we also compare, for all experimental settings with S/D [5]. S/D uses the ratio of the time series values at end-systole and diastole, and labels baseline a time series when this ratio is over 1 and ischemic otherwise. In a synthetic dataset, we assume the position of diastole as the last point in the time series and estimate ES as the location in the utilized kernel (or shifted kernel) of the sub-model that generated the data with the maximum intensity value. ICA outperforms S/D in all cases and exhibits less variability.

To evaluate robustness to additive noise, we varied σ_{ϵ} of the noise model $\mathcal{N}(\mu_{\epsilon}, \sigma_{\epsilon})$ considering $\hat{N} = 600$ and $\hat{N}_i = 200$. In-

creasing σ_{ϵ} from 0 to 16 results in a 6% drop in the performance of the ICA (92% to 86%) while the drop of S/D is significantly higher, 11% (81% to 70%).

To put the obtained accuracy of ICA in perspective, when $\hat{N}_i = 100$ and $\hat{N} = 300$, its accuracy is $88\pm5\%$. This value is close to the one obtained with the supervised method of SVM on real data which was used for validation (Table II). Although here we used synthetic data, it illustrates that unsupervised identification of ischemia on the basis of time series extracted from myocardial pixels on the basis of CP-BOLD MRI is possible, and that fMRI literature and methods, particularly those based on decompositions, can provide fruitful inspiration.

E. Investigation of Shifts With Circulant Only MC-DLA

We previously saw that a characteristic CP-BOLD curve exists and that shifted versions of it are utilized throughout the dataset (Fig. 4). Such shifts may arise due to variations in endsystole (ES) location as discussed previously. In addition, the presence of shifts across different territories of the same subject, is likely attributed to physiological differences among different arteries [5], [30]. Relying on the baseline dataset and MC-DLA(28, 0, 1, 0) we can construct a statistical procedure to test for their existence. Based on the learned circulant kernel $c \in \mathbb{R}^n$ and MC-DLA(28, 0, 1, 0) we can obtain the importance of each shift for each subject k, denoted as $I_k(i)$. If no shifts exist between time series of the same subject, $I_k(i)$ should be dominated by a single high valued peak; thus, less collective importance on the remaining peaks is expected and the lower their collective importance the less shifts would exist. Using this simple observation, on average we find $72 \pm 6\%$ importance outside the main peak, and when tested versus 25% the null was rejected (p-value $< 10^{-4}$).

IV. DISCUSSION AND CONCLUSION

The overarching goals of this work are 1) to generate synthetic CP-BOLD time series with good resemblance to real data; and 2) to provide insight into the characteristics of the differential variations in BOLD signal intensities.

To achieve these goals, the proposed composite dictionary learning algorithm provides a proper, data-driven, representation of the input time series. Each time series is represented by a linear combination between one circulant atom (the kernel or one of its shifts), and other dictionary atoms. We are able to identify a characteristic behavior under baseline conditions which supports the notion that the BOLD signal is cardiac phase dependent and that it can vary across myocardial territories, possibly due to physiological/anatomical differences between subjects [5], [30]. Most importantly, our results suggest that all baseline CP-BOLD time series should follow this behavior, which we refer to as the *characteristic CP-BOLD curve*. These findings, demonstrated here for the first time are remarkable and provide further insight into cardiac physiology.

The availability of the proposed synthetic generators allows us to build and test detection algorithms independently of accurate myocardial segmentation and registration and without requiring new experimental data. In this paper, as proof-of-concept, with synthetic data we show that classical unsupervised methods (based on ICA) can potentially be used to identify areas of ischemia in CP-BOLD MRI, outperforming previously considered methods [5]. Highlighting the importance for precision in myocardial segmentation and registration, we also show that the number of available time series and their quality (i.e., noise level) directly affects the accuracy of ischemia detection. In the future, with our generator, for example, we could also test (and quantify) the impact of shift variability on the accuracy and design of ischemia detection algorithms. We anticipate a lower accuracy if the presence of physiological shifts are not accounted for, when designing the algorithm.

Furthermore, insights on how the BOLD effect varies according to cardiac phase and disease status, may be integrated within image-based generators (e.g., [3]), to create synthetic but realistic imaging stacks. These stacks can be used to facilitate the development of segmentation and registration methods tailored to cardiac BOLD imaging. Similarly these insights can improve the design of multi-state hemoglobin relaxation models. These models can be used to examine via simulation how MR sequence design and imaging protocols affect detection methods down the line.

Being the first study to propose synthetic generation of CP-BOLD time series, this study could not compare with any prior method in the literature, which can be seen as a limitation. Nevertheless, for a fair comparison we included a PCA-derived decomposition for modeling which follows similar principles to the one proposed here. However, their comparison beyond validation accuracy must include the capability to extract the object of interest, i.e., the characteristic CP-BOLD curve. The proposed method cannot directly model spatial correlations (i.e., that time series extracted from neighboring territories are correlated), but postgeneration time series can be smoothed to introduce correlation. Clearly, we need a sufficient number of training data with ground truth to learn the two DBMA models. For baseline conditions this is much easier than ischemic conditions, due to limitations in spatial resolution, potential partial volume effects and, in this study, the reliance on LGE images to identify the area at risk. Further improvements in MRI techniques that increase spatial resolution and fidelity could alleviate such issues in the future. In addition, adopting techniques that could register information of area at risk with cine (BOLD) MRI [50], [51] could improve the identification of ischemic time series for training and validation purposes.

From a technical perspective, the newly proposed MC-DLA is able to provide a unique decomposition of input time series as linear combinations between structured (circulant) and unstructured (generic) dictionary components or atoms. Learning this model is done efficiently using a newly proposed circulant [40] and a well-known dictionary learning mechanisms. Based on this decomposition we statistically model all low dimensional subspaces found. Outside the context of cardiac MRI, MC-DLA could also decompose fMRI data, particularly if shift-invariant decomposition is sought-after [52]–[54]. Or, it can be used more generally for modeling and clustering data [48] or imaging applications [55]. We also provide a fitness measure to identify the internal dimensions of the model such that representation capabilities and model complexity are well balanced. This is necessary since the traditional dictio-

nary learning framework does not naturally accommodate for the compact dictionary and subspace structure imposed in our models. The proposed fitness measure can be applied whenever dictionary and subspace models with specific structure are applied [48].

In conclusion, we show that the proposed method can generate realistic CP-BOLD time series. It learns and represents the data efficiently, taking advantage of a unique sparse decomposition framework to produce parsimonious models that are easy to use. Its application towards the development of better ischemia detection algorithms and also the simulation of imaging scenarios when designing CP-BOLD acquisitions accelerates the clinical translation of this truly noninvasive and repeatable method for the assessment of ischemia.

REFERENCES

- W. P. Segars, D. S. Lalush, and B. M. W. Tsui, "A realistic spline-based dynamic heart phantom," *IEEE Trans. Nucl. Sci.*, vol. 46, no. 3, pp. 503–506, Jun. 1999.
- [2] J. Montagnat and H. Delingette, "4D deformable models with temporal constraints: Application to 4D cardiac image segmentation," *Med. Image Anal.*, vol. 9, pp. 87–100, 2005.
- [3] A. Prakosa, M. Sermesant, H. Delingette, S. Marchesseau, E. Saloux, P. Allain, N. Villain, and N. Ayache, "Generation of synthetic but visually realistic time series of cardiac images combining a biophysical model and clinical images," *IEEE Trans. Med. Imag.*, vol. 32, no. 1, pp. 99–109, Jan. 2013.
- [4] C. F. Beckmann and S. M. Smith, "Probabilistic independent component analysis for functional magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 23, no. 2, pp. 137–152, Feb. 2004.
- [5] S. A. Tsaftaris, X. Zhou, R. Tang, D. Li, and R. Dharmakumar, "Detecting myocardial ischemia at rest with cardiac phase-resolved blood oxygen level-dependent cardiovascular magnetic resonance," *Circ. Cardiovasc. Imag.*, vol. 6, no. 2, pp. 311–319, 2013.
- [6] R. Dharmakumar, J. Arumana, R. Tang, K. Harris, Z. Zhang, and D. Li, "Assessment of regional myocardial oxygenation changes in the presence of coronary artery stenosis with balanced SSFP imaging at 3.0T: Theory and experimental evaluation in canines," *J. Magn. Reson. Imag.*, vol. 27, no. 5, pp. 1037–1045, 2008.
- [7] X. Zhou, S. A. Tsaftaris, Y. Liu, R. Tang, R. Klein, S. Zuehlsdorff, D. Li, and R. Dharmakumar, "Artifact-reduced two-dimensional cine steady state free precession for myocardial blood-oxygen-level-dependent imaging," *J. Magn. Reson. Imag.*, vol. 31, no. 4, pp. 863–871, 2010.
- [8] T. Marwick, B. Willemart, A. M. D'Hondt, T. Baudhuin, W. Wijns, J. M. Detry, and J. Melin, "Selection of the optimal nonexercise stress for the evaluation of ischemic regional myocardial dysfunction and malperfusion. Comparison of dobutamine and adenosine using echocardiography and 99 MTC-MIBI single photon emission computed tomography," *Circulation*, vol. 87, no. 2, pp. 345–354, 1993.
- [9] U. S. Food Drug Admin., FDA warns of rare but serious risk of heart attack and death with cardiac nuclear stress test drugs Lexiscan (regadenoson) and Adenoscan (adenosine), 2013 [Online]. Available: http://www.fda.gov/Drugs/DrugSafety/ucm375654.htm
- [10] S. A. Tsaftaris, R. Tang, X. Zhou, D. Li, and R. Dharmakumar, "Ischemic extent as a biomarker for characterizing severity of coronary artery stenosis with blood oxygen-sensitive MRI," *J. Magn. Reson. Imag.*, vol. 35, no. 6, pp. 1338–1348, 2012.
- [11] M.-P. Jolly, "Automatic segmentation of the left ventricle in cardiac MR and CT images," *Int. J. Comput. Vis.*, vol. 70, no. 2, pp. 151–163, 2006.
- [12] W. Feng, H. Nagaraj, H. Gupta, S. G. Lloyd, I. Aban, G. J. Perry, D. A. Calhoun, L. J. Dell'Italia, and T. S. Denney, "A dual propagation contours technique for semi-automated assessment of systolic and diastolic cardiac function by CMR," *J. Cardiovas. Magn. Reson.*, vol. 11, pp. 30–43, 2009.
- [13] C. Petitjean and J.-N. Dacher, "A review of segmentation methods in short axis cardiac MR images," *Med. Image Anal.*, vol. 15, no. 2, pp. 169–184, 2011.

- [14] A. Suinesiaputra, B. R. Cowan, A. O. Al-Agamy, M. A. AlAttar, N. Ayache, A. S. Fahmy, A. M. Khalifa, P. Medrano-Gracia, M.-P. Jolly, A. H. Kadish, D. C. Lee, J. Margeta, S. K. Warfield, and A. A. Young, "A collaborative resource to build consensus for automated left ventricular segmentation of cardiac MR images," *Med. Image Anal.*, vol. 18, no. 1, pp. 50–62, 2014.
- [15] C. Feng, C. Li, D. Zhao, C. Davatzikos, and H. Litt, "Segmentation of the left ventricle using distance regularized two-layer level set approach," in *Proc. MICCAI*, 2013, pp. 477–484.
- [16] W. Bai, W. Shi, D. P. O'Regan, T. Tong, H. Wang, S. Jamil-Copley, N. S. Peters, and D. Rueckert, "A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: Application to cardiac MR images," *IEEE Trans. Med. Imag.*, vol. 32, no. 7, pp. 1302–1315, Jul. 2013.
- [17] T. Makela, P. Clarysse, O. Sipila, N. Pauna, Q. Pham, T. Katila, and I. E. Magnin, "A review of cardiac image registration methods," *IEEE Trans. Med. Imag.*, vol. 21, no. 9, pp. 1011–1021, Sep. 2002.
- [18] D. Perperidis, R. Mohiaddin, and D. Rueckert, "Spatio-temporal freeform registration of cardiac MR image sequences," *Med. Image Anal.*, vol. 9, no. 5, pp. 441–456, 2005.
- [19] Y. Ou, A. Sotiras, N. Paragios, and C. Davatzikos, "DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting," *Med. Image Anal.*, vol. 15, no. 4, pp. 622–639, 2011.
- [20] Y. Ou, D. H. Ye, K. M. Pohl, and C. Davatzikos, "Validation of DRAMMS among 12 popular methods in cross-subject cardiac MRI registration," in WBIR, 2012, pp. 209–219.
- [21] A. F. Frangi, W. J. Niessen, and M. A. Viergever, "Three-dimensional modeling for functional analysis of cardiac images, a review," *IEEE Trans. Med. Imag.*, vol. 20, no. 1, pp. 2–5, Jan. 2001.
- [22] F. Li, J. A. White, M. Rajchl, A. Goela, and T. M. Peters, "Generation of synthetic 4D cardiac CT images by deformation from cardiac ultrasound," in *Augmented Environments for Computer-Assisted Interventions*, L., Ed. *et al.* New York: Springer, 2013, vol. 7815, pp. 132–141, ser. LNCS.
- [23] R. Maitra and J. Riddles, "Synthetic magnetic resonance imaging revisited," *IEEE Trans. Med. Imag.*, vol. 29, no. 3, pp. 895–902, Mar. 2010.
- [24] S. Aja-Fernandez, L. Cordero-Grande, and C. Albewla-Lopez, "A MRI phantom for cardiac perfusion simulation," in *Proc. 9th IEEE Int. Symp. Biomed. Imag.*, 2012, pp. 638–641.
- [25] A. Chiribiri, A. Schuster, M. Ishida, G. Hautvast, N. Zarinabad, G. Morton, J. Otton, S. Plein, M. Breeuwer, P. Batchelor, T. Schaeffter, and E. Nagel, "Perfusion phantom: An efficient and reproducible method to simulate myocardial first-pass perfusion measurements with cardiovascular magnetic resonance," *Magn. Res. Med.*, vol. 69, no. 3, pp. 698–707, 2013.
- [26] C. Tobon-Gomez, F. M. Sukno, B. H. Bijnens, M. Huguet, and A. F. Frangi, "Realistic simulation of cardiac magnetic resonance studies modeling anatomical variability, trabeculae, and papillary muscles," *Magn. Res. Med.*, vol. 65, no. 1, pp. 280–288, 2011.
- [27] R. Dharmakumar, J. Hong, J. H. Brittain, D. B. Plewes, and G. A. Wright, "Oxygen-sensitive contrast in blood for steady-state free precession imaging," *Magn. Reson. Med.*, vol. 53, no. 3, pp. 574–583, 2005.
- [28] T. D. Karamitsos, L. Leccisotti, J. R. Arnold, A. Recio-Mayoral, P. Bhamra-Ariza, R. K. Howells, N. Searle, M. D. Robson, O. E. Rimoldi, P. G. Camici, S. Neubauer, and J. B. Selvanayagam, "Relationship between regional myocardial oxygenation and perfusion in patients with coronary artery disease: Insights from cardiovascular magnetic resonance and positron emission tomography," *Circ. Cardiovasc. Imag.*, vol. 3, no. 1, pp. 32–40, 2010.
- [29] K. S. McCommis, T. A. Goldstein, D. R. Abendschein, P. Herrero, B. Misselwitz, R. J. Gropler, and J. Zheng, "Quantification of regional myocardial oxygenation by magnetic resonance imaging: Validation with positron emission tomography," *Circ. Cardiovasc. Imag.*, vol. 3, no. 1, pp. 41–46, 2010.
- [30] Y. Ootaki, C. Ootaki, K. Kamohara, M. Akiyama, F. Zahr, J. Kopcak, M. W., R. Dessoffy, and K. Fukamachi, "Phasic coronary blood flow patterns in dogs vs. pigs: An acute ischemic heart study," *Med. Sci. Monit.*, vol. 14, no. 10, pp. 193–197, 2008.
- [31] P. Kellman, A. E. Arai, E. R. McVeigh, and A. H. Aletras, "Phasesensitive inversion recovery for detecting myocardial infarction using gadolinium-delayed hyperenhancement," *Magn. Res. Med.*, vol. 47, no. 2, pp. 372–383, 2002.

- [32] S. A. Tsaftaris, V. Andermatt, A. Schlegel, A. K. Katsaggelos, D. Li, and R. Dharmakumar, "A dynamic programming solution to tracking and elastically matching left ventricular walls in cardiac cine MRI," in *Proc. Int. Conf. Image Proc.*, 2008, pp. 2980–2983.
- [33] D. L. Donoho, "High-dimensional data analysis: The curses and blessings of dimensionality," AMS Math Challenges Lecture, 2000.
- [34] A. Criminisi, J. Shotton, and E. Konukoglu, "Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning," *Found. Trends. Comput. Graph. Vis.*, vol. 7, no. 2–3, pp. 81–227, 2012.
- [35] R. Jenatton, R. Gribonval, and F. Bach, "Local stability and robustness of sparse dictionary learning in the presence of noise," arXiv:1210. 0685, 2012.
- [36] Z. J. Xiang, H. Xu, and P. J. Ramadge, "Learning sparse representations of high dimensional data on large scale dictionaries," in *Adv. Neural Inf. Process. Syst.*, 2011, pp. 900–908.
- [37] J. Silva, M. Chen, Y. C. Eldar, G. Sapiro, and L. Carin, "Blind compressed sensing over a structured union of subspaces," arXiv:1103. 2469, 2011.
- [38] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1982–2001, Apr. 2010.
- [39] G. Pope, C. Aubel, and C. Studer, "Learning phase-invariant dictionaries," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2013, pp. 5979–5983.
- [40] C. Rusu, B. Dumitrescu, and S. A. Tsaftaris, "Explicit shift-invariant dictionary learning," *IEEE Signal Process. Lett.*, vol. 21, no. 1, pp. 6–9, Jan. 2014.
- [41] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [42] C. Rusu and B. Dumitrescu, "Stagewise K-SVD to design efficient dictionaries for sparse representations," *IEEE Signal Process. Lett.*, vol. 19, no. 10, pp. 631–634, Oct. 2012.
- [43] Y. Pati, R. Rezaiifar, and P. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with application to wavelet decomposition," in *Asilomar Conf. Signals, Syst. Comput.*, 1993, vol. 1, pp. 40–44.
- [44] R. Redner and H. Walker, "Mixture densities, maximum likelihood and the EM algorithm," SIAM Rev., vol. 26, no. 2, pp. 195–239, 1984.
- [45] L. Ruan, M. Yuan, and H. Zou, "Regularized parameter estimation in high-dimensional Gaussian mixture models," *Neural Comput.*, vol. 23, no. 6, pp. 1605–1622, 2011.
- [46] R. Gribonval and P. Vandergheynst, "On the exponential convergence of matching pursuits in quasi-incoherent dictionaries," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 255–261, Jan. 2006.
- [47] R. Mazhar and P. D. Gader, "EK-SVD: Optimized dictionary design for sparse representations," in *Proc. 19th Int. Conf. Pattern Recognit.*, 2008, pp. 1–4.
- [48] Y. Ma, A. Y. Yang, H. Derksen, and R. Fossum, "Estimation of subspace arrangements with applications in modeling and segmenting mixed data," *SIAM Rev.*, vol. 50, no. 3, pp. 413–458, 2008.
- [49] C. J. C. Burges, "A tutorial on support vector machines for pattern recognition," *Data Min. Knowl. Discov.*, vol. 2, pp. 121–167, 1998.
- [50] Y. Liu, H. Xue, C. Guetter, M. P. Jolly, N. Chrisochoides, and J. Guehring, "Moving propagation of suspicious myocardial infarction from delayed enhanced cardiac imaging to cine MRI using hybrid image registration," in *Proc. IEEE Int. Symp. Biomed. Imag.: From Nano to Macro*, 2011, pp. 1284–1288.
- [51] L. Cordero-Grande, S. Merino-Caviedes, X. Alba, R. Figueras i Ventura, A. F. Frangi, and C. Alberola-Lopez, "3D fusion of cine and late-enhanced cardiac magnetic resonance images," in *Proc. 2012 9th IEEE Int. Symp. Biomed. Imag.*, 2012, pp. 286–289.
- [52] M. Mørupa, L. K. Hansena, S. M. Arnfredb, L. H. Limc, and K. H. Madsena, "Shift-invariant multilinear decomposition of neuroimaging data," *NeuroImage*, vol. 42, no. 4, pp. 1439–1450, 2008.
- [53] H. Eavani, R. Filipovych, C. Davatzikos, T. D. Satterthwaite, R. E. Gur, and R. C. Gur, "Sparse dictionary learning of resting state fMRI networks," in *Pattern Recognit. NeuroImag.*, 2012, pp. 73–76.
- [54] H. Eavani, T. D. Satterthwaite, R. E. Gur, R. C. Gur, and C. Davatzikos, "Unsupervised learning of functional network dynamics in resting state fMRI," in *Information Processing in Medical Imaging*. New York: Springer, 2013, vol. 7917, Lecture Notes Comput. Sci., pp. 426–437.
- [55] N. Raoy, B. Recht, and R. D. Nowaky, "Signal recovery in unions of subspaces with applications to compressive imaging," arXiv:1209. 3079v1, 2012.