Integrated Segmentation and Interpolation of Sparse Data

Adeline Paiement, Majid Mirmehdi, *Senior Member, IEEE*, Xianghua Xie, *Member, IEEE*, and Mark C. K. Hamilton

***Abstract*— We address the two inherently related problems of segmentation and interpolation of 3D and 4D sparse data and propose a new method to integrate these stages in a level set framework. The interpolation process uses segmentation information rather than pixel intensities for increased robustness and accuracy. The method supports any spatial configurations of sets of 2D slices having arbitrary positions and orientations. We achieve this by introducing a new level set scheme based on the interpolation of the level set function by radial basis functions. The proposed method is validated quantitatively and/or subjectively on artificial data and MRI and CT scans and is compared against the traditional sequential approach, which interpolates the images first, using a state-of-the-art image interpolation method, and then segments the interpolated volume in 3D or 4D. In our experiments, the proposed framework yielded similar segmentation results to the sequential approach but provided a more robust and accurate interpolation. In particular, the interpolation was more satisfactory in cases of large gaps, due to the method taking into account the global shape of the object, and it recovered better topologies at the extremities of the shapes where the objects disappear from the image slices. As a result, the complete integrated framework provided more satisfactory shape reconstructions than the sequential approach.**

***Index Terms*— 3D/4D object modeling, segmentation, interpolation, level set methods, RBF.**

I. Introduction

MODELING from 3D and 4D tomographic volumes raises two closely intertwined issues, segmentation and interpolation. Segmentation is required to partition the 3D space containing the object and to distinguish data points belonging to the object from background points. Typical datasets include sets of 2D slices, ranging from simple stacks of parallel slices to more complicated spatial configurations with slices having various positions and orientations (see Fig. 1). 3D datasets can contain slices of both the object of interest and its environment, and independent 2D segmentation of the constituent 2D slices is not suitable to reconstruct 3D or 4D volumes, since these segmentations may not be consistent with each other. Therefore, all the slices are better segmented simultaneously in 3D or 4D. Interpolation is necessary since data often do not span the whole 3D space – thus offering only partial support to segmentation. Some medical imaging modalities require integrating the signal over a thick slice of the volume in order to improve the signal quality by increasing the signal to noise ratio. This sets an undesirable limit on the possible slice spacings. For example, for cine cardiac MRIs produced by a 1.5T scanner, a typical slice thickness is 7mm, hence slice spacing is usually 7mm or larger. In addition, clinicians sometimes choose to acquire largely spaced slices in order to decrease the acquisition time and reduce patient discomfort, resulting in very sparse volumes containing large gaps between the 2D slices, as illustrated in Fig. 1. This is a common practice, especially for imaging modalities such as cardiac MRI, when image capture can be very difficult for patients who are asked to remain perfectly still until the end of the acquisition and to hold their breath repeatedly for around 10 seconds and occasionally up to 30 seconds at a time. Such data acquisition protocols result in gaps with widths of around 8mm to 16mm for in-plane pixel sizes of $\sim$1.8mm. These gaps make the 3D segmentation even more complicated and data must be interpolated in order to reconstruct a more representative 3D volume. In the 4D case, data must also be interpolated between the available time frames. Clearly, the success of one stage (segmentation or interpolation) depends on the accuracy of the other.

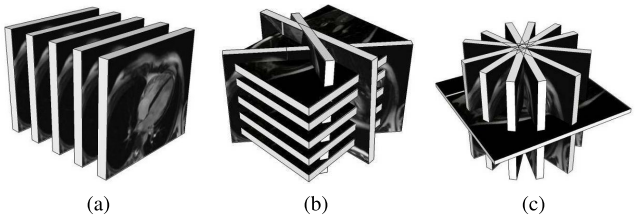

(a) (b) (c)

Fig. 1. Examples of sparse sets of slices – slice thickness is imposed by scanner's limitations, while spacing is adjusted as a compromise between accuracy and patient comfort. (a) simple stack of parallel slices, (b) standard spatial configuration of a cardiac MRI: stack of parallel short-axis slices plus a few long-axis ones, and (c) radial dataset.

Traditionally, segmentation and interpolation are performed iteratively. Two sequential approaches exist which perform these two stages in opposing order. Some works first segment the slices independently, and then interpolate a surface from the 2D contours, e.g. [1], [2]. As mentioned earlier,

Manuscript received February 8, 2013; revised July 16, 2013 and September 18, 2013; accepted September 23, 2013. Date of publication October 23, 2013; date of current version November 7, 2013. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Richard J. Radke.

A. Paiement and M. Mirmehdi are with the Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K. (e-mail: csatmp@bristol.ac.uk; majid@cs.bris.ac.uk).

X. Xie is with the Department of Computer Science, University of Swansea, Swansea SA2 8PP, U.K. (e-mail: x.xie@swansea.ac.uk).

M. C. K. Hamilton is with the National Institute for Health Research, Cardiovascular Biomedical Research Unit and the Bristol Heart Institute, Bristol Royal Infirmary, Bristol BS2 8HW, U.K. (e-mail: mark.hamilton@uhbristol.nhs.uk).

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/TIP.2013.2286903

TABLE I
SUMMARY OF EXISTING APPROACHES FOR 3D MODELING FROM SPARSE MEDICAL DATA

Approach		Type of segmentation	Type of interpolation	Main drawbacks	
Sequential	Segmentation followed by interpolation	2D	Surface interpolation	Robustness of segmentation	First step
	Interpolation followed by segmentation	3D	Image interpolation	Ill-posed interpolation problem	is critical
Model based		3D, registration and deformation of a model	Model continuity	Training + Lack of flexibility	
Interpolated level set surface [8]		3D level set	Diffusion of the level set surface in gaps	Robustness of the interpolation	

the drawback of this approach is that the independent 2D segmentations may not be consistent with each other, and incorrect segmentations would not provide a reliable ground for the shape interpolation. Liu *et al.* proposed in [3] a new method to interpolate surfaces from segmentation contours obtained from image slices having arbitrary positions and orientations. In their work, the segmentations are performed manually to ensure they are correct. More commonly, other works perform the interpolation of the slices first in order to reconstruct a 3D volume, e.g. as in [4]–[7], and then perform the segmentation. The problem of interpolation between images having different orientations is ill-posed, and most of these works, e.g. [7], apply only to one or two stacks of parallel and equally spaced images. In addition, methods based on the interpolation of image intensities across and in-between frames may be very sensitive to differences of gain and contrast in the images, thus producing interpolation artifacts which may bias the segmentation stage that follows. This deficiency was demonstrated in our earlier work [8], where interpolation produced dark lines which attract the segmenting contour. In order to avoid creating interpolation artifacts, Woo *et al.* proposed in their work on superresolution by volume fusion [9] to equalize the intensities of corresponding regions in three non-isotropic volumes using an intensity matching method based on spline regression, prior to reconstructing a high resolution fused volume. However modifying the intensities of medical images may be controversial because of the risk in erasing valuable information and creating erroneous features.

These sequential approaches lack robustness and an underperformance of one stage can lead to further challenges in, and even failure of, the subsequent stage. However, only a few attempts have been made to combine these stages in a single framework. Some works have segmented sparse volumes made up of 2D slices by registering and deforming a model on the images, thus relying on the continuity of the model to handle the gaps, e.g. [10] with ASM models, [11] with combined 2D AAMs and a 3D shape model, [12] with a triangulated surface model, and [13] with a parametric model. This can be seen as integrated segmentation of the images and interpolation of the shape of the segmented object in the gaps. These methods have the drawbacks inherent in all model based approaches – they require prior knowledge and a training phase, and lack flexibility. Another method which combines segmentation and interpolation was presented in our earlier work in [8]. It is based on the flexible level set method and does not require prior knowledge. We will describe the method in more detail in the next section. Table I presents a summary of existing methods for object modeling from 3D sparse medical data, together with their main drawbacks.

The novel contributions of this work can be summarized as follows. We integrate segmentation and interpolation into a new radial basis function (RBF) interpolated level set framework which benefits from the simplicity and flexibility of level set methods, the numerical stability of RBF interpolated level set segmentation methods, and the interpolation abilities of RBFs. To the best of our knowledge, this is the first framework which uses RBFs for both segmentation and interpolation. Furthermore, like in our earlier work [8], the proposed framework interpolates the level set segmenting surface rather than the image intensities, but using a different interpolation scheme which results in better quality interpolations (this will be quantitatively shown in the Results section). This approach is an important aspect of our method since it enables us to handle and exploit images having different gains and contrasts, and even images from different imaging modalities, simultaneously. Further, it uses shape information provided by the segmenting surface in order to yield a more robust interpolation than intensity-based methods. Also, our proposed method is very general and can model objects (i) of any shape, since it is not limited to a given class of objects like model based methods, (ii) from data having any number of 2D slices and any positions and orientations of its slices, (iii) from any image modality by a proper choice of the segmentation method, e.g. edges, intensity or prior knowledge, and (iv) using any strictly positive definite (SPD) RBF.

Next, we review some related works and identify improvements we make on existing methods. Then in Section III, we describe our proposed method and give implementation guidelines,[1] including parameter setting. Validation tests and comparative results are presented in Section IV, and Section V concludes the paper.

II. RELATED WORKS

Object modeling from sparse medical data using a level set method has been carried out, to some extent, in our earlier work on segmentation and interpolation of sets of slices [8]. Level set methods evolve a contour or a surface driven

[1]Our software will be publicly available at http://www.bristol.ac.uk/vi-lab/projects/sparseobjectmodelling/.

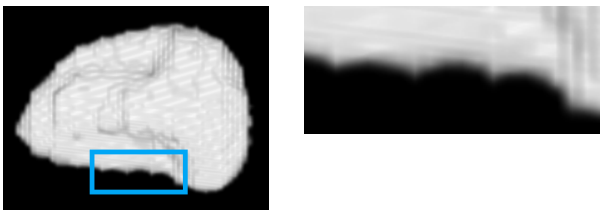

(a) (b)

Fig. 2. Interpolation by surface diffusion in gaps using an MCF – the modeled object is the left ventricle cavity of a heart reconstructed from a MRI dataset. (a) global view, (b) zoom of the highlighted part of (a).

typically by a data term which matches the contour or surface to the data, and a smoothing term which ensures the contour or surface remains smooth. In [8], the data term evolves the parts of the level set implicit surface which are in the planes of the 2D slices in order to segment those images, and the implicit surface is diffused in the gaps between the slices using a Mean Curvature Flow (MCF) [14]. This curvature based smoothing term ensures continuity and smoothness in the gaps between the 2D slices. Any segmentation algorithm may be used to segment the 2D slices, and it is even possible to use different algorithms simultaneously to process images having different gains and contrasts, and images from different modalities acquired in different scans, provided that they are properly aligned and scaled. A fundamental shortcoming of MCF is that it tends to produce minimal surfaces which, depending on the application, may not be appropriate. Indeed, Chopp and Sethian in [14] showed that a cylinder evolved under the MCF between two fixed rings becomes a catenoid if the rings are close enough, and disappears if the distance between the rings is larger than a maximum distance. This caused the numerous catenoid-like shapes in Fig. 2. It may be more suitable to take into account the global geometry of the object in order to produce a better interpolation. Surfaces yielded by diffusion using the Surface Diffusion Flow (SDF) [15] are minimal curvature surfaces and tend to be more visually satisfactory, but similarly to MCF, SDF is also not easily scalable. In particular, for large datasets and/or large gaps between the slices, many SDF iterations are required at each segmentation iteration in order to successfully drive the level set implicit surface from one side of the gap to the other. In addition, for bigger volumes or gaps, the method often fails to propagate the implicit surface satisfactorily, as in Figs. 3a and 3b where it is not interpolated smoothly in the gaps of a vertebra dataset, and in Fig. 3c where it failed to propagate to the upper and lower parts of brain ventricles, which extend further up and down inside the skull, as shown by the raw data of Figs. 3d and 3e. Our proposed method performs interpolation of the level set implicit surface in a novel way which exploits interpolation properties of RBFs and yields a more satisfactory interpolation of the shape of the object than [8]. In particular, [8] uses a local approach based on diffusion of the level set implicit surface in the gaps, while the proposed method can achieve a global interpolation of the shape.

RBFs have been applied to both segmentation and interpolation in the past, as shown in Table II, and are particularly

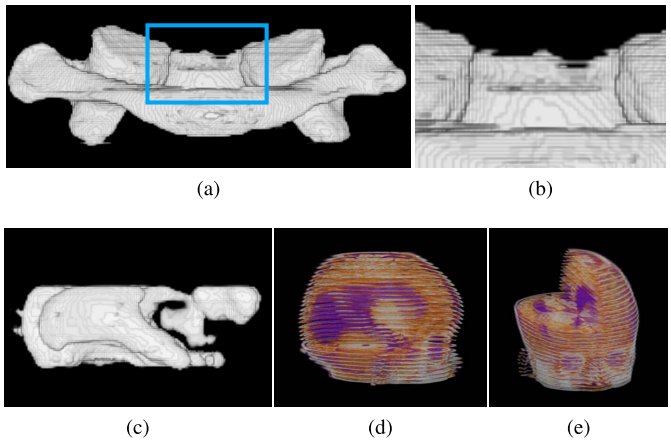

(a) (b) (c) (d) (e)

Fig. 3. Interpolation by surface diffusion in gaps using an SDF. Vertebra from a CT-scan: (a) view from the front, and (b) zoom of the highlighted part of (a), and (c) brain ventricles from an MRI dataset, viewed from the side. Both datasets are made up of parallel slices. The brain dataset is shown in (d) from the side and (e) from the front, with parts removed in order to allow a better visualization of the ventricles, in purple.

TABLE II

SUMMARY OF THE MAIN EXISTING USES OF RBFS

Use of RBF	Segmentation	Interpolation
Interpolation of scattered data, e.g. [16]		×
Object modeling from point clouds, e.g. [17]		×
Interpolation of a level set implicit function, e.g. [18]–[20]	×	
Proposed method	×	×

popular for interpolation of scattered data, e.g. [16]. In [17], Carr *et al.* used this interpolation to fit an implicit surface to scattered 3D points in order to model objects from 3D point clouds. Recently, Wang *et al.* [18] presented a new use of RBFs by segmenting images using a level set function interpolated by RBFs, thus transforming the partial differential equation (PDE) governing the motion of the level set's interface into a simpler ordinary differential equation (ODE) and obtaining a more numerically stable scheme. However, this new level set segmentation method requires inverting and storing a matrix whose size depends on the size of the dataset. It is therefore unsuitable for processing large datasets such as 3D and 4D data. This work was then further improved by some authors who addressed this inversion issue by deriving new ODEs using a variational method, resulting in more memory and computationally efficient schemes. Amongst them, Slabaugh *et al.* in [19] proposed to use any RBFs which are anisotropic with an analytic derivative, such as 2D anisotropic Gaussians, and derived an evolution equation for both region-based and edge-based segmentation algorithms. The number, locations and anisotropic properties[2] of the RBFs are updated at each iteration in addition to their usual weighting coefficients, and some RBFs may be merged and added when required. In comparison, our new RBF interpolated level set framework is more easily implemented since it neither requires

[2] σ_x and σ_y for the 2D anisotropic Gaussians.

updating the properties of each individual RBF, but only their weights, nor the merging and adding of RBFs. Bernard *et al.* in [20] used B-spline RBFs to interpolate the level set function, and derived an evolution equation for it, for the case of region-based segmentation only.

In summary, RBF based approximation methods do not offer any segmentation of the data or of the volume containing them, while methods such as [19], [20] do not offer interpolation and can not be used on sparse data. Our proposed method combines these two uses of RBFs – segmentation and interpolation – in order to make sense of scattered data. Our framework benefits from the high numerical stability of RBF interpolated level set segmentation methods, so that large time steps can be used, and frequent renormalizations to maintain a signed distance function are not required, like they are for conventional level sets. It is a variational method and is therefore more computationally efficient than methods such as [18], so it can handle large datasets such as 3D and 4D data. Further, as mentioned earlier, our proposed framework is more general than methods such as e.g. [19], [20], since any segmentation criteria and any SPD RBF may be used, in any number of dimensions.

III. Proposed Integrated Segmentation and Interpolation Framework

Let $\phi(\mathbf{x})$ be a level set function and $\psi_i(\mathbf{x}) = \psi_i(\|\mathbf{x}\|)$ a SPD RBF centered on a control point $\mathbf{x}_i$. ϕ may be interpolated as

$$\phi(\mathbf{x}) = \sum_{i=1}^{N} \alpha_i \psi_i(\mathbf{x}) = \sum_{i=1}^{N} \alpha_i \psi(\mathbf{x} - \mathbf{x}_i), \tag{1}$$

where α_i are weighting coefficients. The method proposed by [18] to evolve ϕ through $\alpha = [\alpha_1 \dots \alpha_N]^T$ involves inverting and storing an $N \times N$ matrix, which may be very time consuming and memory hungry when dealing with the large number of points N which compose a 3D or 4D dataset. Instead, we derive the evolution of α by minimizing an energy functional $E[\phi]$ governing the segmentation of the space Ω:

$$E[\phi] = \int_{\Omega} F[\phi(\mathbf{x})]\, d\mathbf{x} = \int_{\Omega} F\left[\sum_{i=1}^{N} \alpha_i \psi_i(\mathbf{x})\right] d\mathbf{x}. \tag{2}$$

$F[\phi]$ may be any functional and is defined by the chosen segmentation method. Conventional variational level set methods derive an evolution equation for ϕ from (2) through a gradient descent method as

$$\frac{\partial \phi}{\partial t}(\mathbf{x}) = -\frac{\partial E}{\partial \phi}(\mathbf{x}) = -\frac{\partial F}{\partial \phi}(\mathbf{x}). \tag{3}$$

Using (1), (3), and the chain rule, a gradient descent method on (2) yields the evolution equation for α_i:

$$\begin{aligned} \frac{d\alpha_i}{dt} &= -\frac{\partial E}{\partial \alpha_i} \\ &= -\int_{\Omega} \frac{\partial F}{\partial \alpha_i}(\mathbf{x})\, d\mathbf{x} \\ &= -\int_{\Omega} \frac{\partial F}{\partial \phi}(\mathbf{x}) \frac{\partial \phi}{\partial \alpha_i}(\mathbf{x})\, d\mathbf{x} \\ &= \int_{\Omega} \frac{\partial \phi}{\partial t}(\mathbf{x})\, \psi_i(\mathbf{x})\, d\mathbf{x}. \end{aligned} \tag{4}$$

Let us rename $\frac{\partial \phi}{\partial t}(\mathbf{x})$ in (3) as $S(\mathbf{x})$. S is the speed of the moving front and is generally defined on the contour C only. Therefore, we can simplify (4) into

$$\frac{d\alpha_i}{dt} = \int_C S(\mathbf{x})\, \psi_i(\mathbf{x})\, d\mathbf{x} \approx \int_{\Omega} \delta_{\varepsilon}(\phi(\mathbf{x}))\, S(\mathbf{x})\, \psi_i(\mathbf{x})\, d\mathbf{x}, \tag{5}$$

where δ_{ε} is an approximation of the Dirac function δ. The restriction of S to the contour C is necessary to ensure the stability of the method, i.e. if the object to be segmented is small with respect to the size of the image, the speeds of the background points would have an overwhelming effect in (4) – they would therefore occlude the influence of the more important speeds of the points located on the object edges. Our experiments showed that δ_{ε} should be larger with increasing flatness of the RBF in order to allow the contour to converge smoothly without oscillating around the edges of the segmented object. In practice, we used

$$\delta_{\varepsilon}(\mathbf{x}) = \begin{cases} \frac{1+\cos\left(\frac{\pi \mathbf{x}}{\varepsilon}\right)}{2\varepsilon} & x \le \varepsilon \\ 0 & x > \varepsilon \end{cases}, \tag{6}$$

with $\varepsilon = 1$ for sharp RBFs and $\varepsilon = 3$ for flatter RBFs. S depends on the chosen segmentation method which may be any method of choice. When segmenting sparse data, no data is available to compute a front's speed in the gaps, so S is set to 0 there.

RBF based interpolation methods usually define one control point per data point. Instead, we choose to define one control point per voxel of a discrete space, allowing (5) to be re-written as a convolution which can then be solved efficiently using the FFT:

$$\begin{aligned} \frac{d\alpha_i}{dt} &= \sum_{C} S(\mathbf{x})\, \psi_i(\mathbf{x}) \\ &\approx \sum_{\Omega} \delta_{\epsilon}(\phi(\mathbf{x}))\, S(\mathbf{x})\, \psi(\mathbf{x}_i - \mathbf{x}) \\ &\approx \left((\delta_{\epsilon}(\phi) \cdot S) \star \psi\right)(\mathbf{x}_i). \end{aligned} \tag{7}$$

$\frac{d\alpha_i}{dt}$ is then used to update α:

$$\alpha^{t+1} = \alpha^t + dt \cdot \frac{d\alpha}{dt}. \tag{8}$$

This step implicitly updates the level set function since ϕ is defined completely by α according to (1). In addition, (1) implies that:

$$\frac{\partial \phi}{\partial t}(\mathbf{x}) = \sum_{i=1}^{N} \frac{d\alpha_i}{dt} \psi_i(\mathbf{x}). \tag{9}$$

Note that this $\frac{\partial \phi}{\partial t}(\mathbf{x})$ is different from the one calculated by conventional level set methods in (3). The initial α may be easily computed in the Fourier domain:

$$\phi(\mathbf{x}) = \sum_{i=1}^{N} \alpha_i \psi(\mathbf{x} - \mathbf{x}_i) = (\alpha \star \psi)(\mathbf{x}), \tag{10}$$

$$\hat{\alpha}(\mathbf{f}) = \frac{\hat{\phi}}{\hat{\psi}}(\mathbf{f}). \tag{11}$$

where $\hat{\alpha}$, $\hat{\phi}$ and $\hat{\psi}$ are the Fourier transforms of α, ϕ and ψ respectively.

1: Initialize a contour C through a level set function $\phi(\mathbf{x})$
2: Compute the initial coefficients α using (10)
3: **repeat**
4: **for all** $\mathbf{x}$ in C **do**
5: Compute $S(\mathbf{x})$ according to the chosen segmentation method
6: **end for**
7: Compute $\frac{d\alpha}{dt}$ using (7)
8: Update α^{t+1} using (8)
9: **until** convergence

Fig. 4. Algorithm 1: Integrated segmentation and interpolation of sparse data.

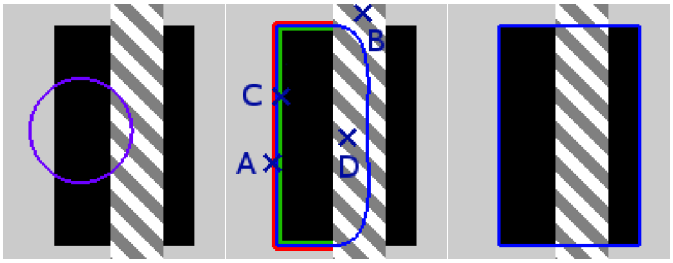

(a) (b) (c)

Fig. 5. Segmentation and interpolation of an image with a gap. (a) initial contour in blue, (b) front speed S – green: positive; red: negative, (c) final segmentation. Note that using (7) and (9) one can assess that $\frac{\partial\phi}{\partial t}$ is respectively negative and positive at points A and C, thus pushing the contour towards the edges of the object, positive at point D, pushing the contour towards the other side of the gap, and negative at point B, preventing it from leaking elsewhere in the gap.

We also note that unlike conventional level set methods, the proposed evolution scheme for the RBF interpolated level set does not use $\nabla\phi$ (see (7) and (8)). In addition, our experiments showed that ϕ tends to remain a smooth function as a result of the smoothing effect of the two convolutions by a decreasing RBF in (7) and (9). It is therefore unnecessary to regularly normalize ϕ to maintain $|\nabla\phi| = 1$. Finally, it is not necessary to compute an extended velocity for points which are not on the contour, as must be done with conventional level sets, because only contour points are used in (7).

Algorithm 1, in Fig. 4, sums up this integrated segmentation and interpolation method, and Fig. 5 illustrates the concept with an example of how an image with a gap can be segmented. It should be stressed that Algorithm 1 can be applied to any sparse medical dataset since it does not make any assumption on the type of data and on the positions and orientations of its 2D slices. The distinction in the processing of different imaging modalities only arises in the computation of the front speed S, which depends on the chosen segmentation algorithm only. This choice is based on the type of data and the imaging modality, and will be discussed in the next section.

A. Choice of RBF

An evolution scheme similar to that presented above can be derived for conditionally positive definite RBF, with the level set function approximated as

$$\phi(\mathbf{x}) = p(\mathbf{x}) + \sum_{i=1}^{N} \alpha_i \psi_i(\mathbf{x}), \tag{12}$$

where p is a polynomial which accounts for the linear and constant portions of ϕ. However, in our experiments such functions yielded rather poor results, with over-smoothed $\frac{\partial\phi}{\partial t}$ and final ϕ. Thus, only globally defined SPD RBFs would be recommended for applications similar to ours.

We wish to emphasize on the importance of the global support of the RBF. Indeed, even if the segmentation stage could work with both compactly supported and globally defined RBFs, the interpolation stage may fail where gaps are larger than their support. We tested Gaussian RBFs and inverse multiquadric RBFs of the form $\psi(\mathbf{x}) = \left(\|\mathbf{x}\|^2 + \gamma^2\right)^{-\frac{\beta}{2}}$, where γ defines the shape of the RBF, and found that the flatness of the RBF has much more influence on the segmentation and interpolation quality than its type. In addition, the fast decay of the Gaussian function makes the evolution of the contour very slow in gaps (see (7)). Therefore, we chose to use inverse multiquadric RBFs, with β equal to the number of dimensions of the data to ensure the decrease is not too steep. The choice of γ allows a finer tuning of the decrease rate.

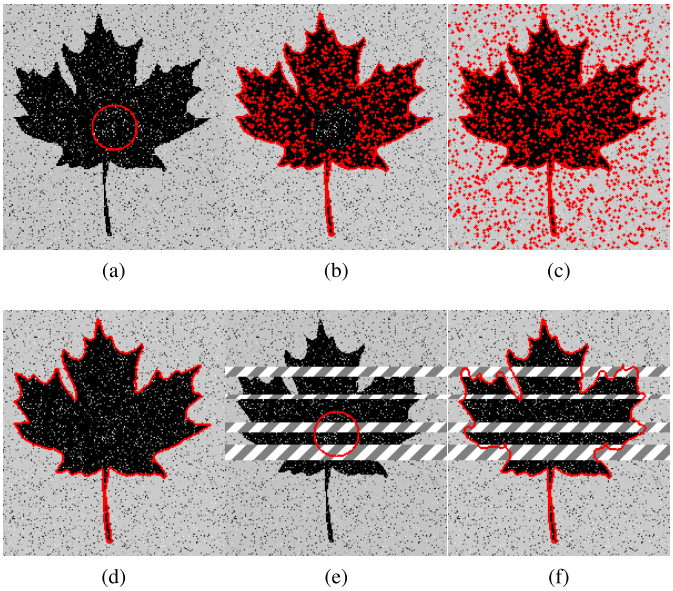

(a) (b) (c) (d) (e) (f)

Fig. 6. Denoising effect: (a) initial noisy image, segmentation by the piecewise constant model used in [23] using (b) conventional narrow-band level set, (c) Chan-Vese level set scheme [23], and (d) proposed level set scheme. (e) initial noisy image with 'slices' of missing information (dashed areas), (f) same as (d) but with gaps being interpolated by the proposed scheme. The piecewise constant model was used to compute $\frac{\partial\phi}{\partial t}$ for (b) and (c) and S for (d) and (f). The schemes of (b) and (c) could not be used with the incomplete image (e) because they can not handle missing information.

B. Choice of Parameter γ

The flatness of ψ, through the value of γ, influences the quality of the segmentation and of the interpolation, as will be discussed below. Its choice depends on the nature of the data, however the same value for γ can be used

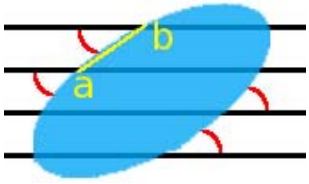

Fig. 7. Example of dataset whose slices (black lines, seen from the side) have a small angle (red) with the surface of the modeled object (blue), resulting in a larger distance ab (yellow) between two neighboring parts of the surface, hence requiring a high degree of smoothing in the gaps.

for processing similar datasets, e.g. in our Experiment 3 in Section IV-D.

The proposed segmentation and interpolation framework has the interesting property of being very robust to noise, as illustrated in Fig. 6. This is due to the two consecutive Gaussian-like smoothings that (7) and (9) apply to S in order to obtain $\frac{\partial \phi}{\partial t}$, resulting in an isolated noisy point and its associated noisy value for S having a reduced influence on the final value of $\frac{\partial \phi}{\partial t}$. Note that these smoothings do not prevent the contour from segmenting sharp corners in Figs. 6d and 6f. The amount of denoising can be adjusted by setting the flatness of ψ through the choice of γ.

A value for γ that is too low would result in an RBF that is too sharp and provides only local interpolation, as only the directly adjacent slices would have an influence on a gap, resulting for example in a 'staircase-like' shape in the case of parallel slices (Fig. 10d). Such situations arise for datasets containing large gaps, and/or in slices that are at a small angle with the object's surface. This latter scenario is illustrated in Fig. 7 and evaluated in experiments on kidney and heart CT-scans reported in Subsection IV-C. In those two cases, a large γ needed to be used in order to obtain a smooth and globally interpolated shape.

On the other hand, a value for γ that is too high, together with inadequate initial conditions, may prevent small holes and/or small objects from being segmented as they would be treated as noise – see examples in Fig. 8k and Fig. 10c. In addition, if the surface of the modeled object has some fine details, such as bumps which are small in comparison to the gap sizes, then a particularly high value of γ results in a highly smoothing segmenting contour that misses the fine details, as illustrated in Fig. 9b.

Hence, the value of γ is highly dependent on the nature of the data, i.e. on the amount of denoising required, the size of the object and of the gaps, and the level of detail on the surface of the object – these can not be automatically evaluated before the object is modeled and, at this point, can only be assessed subjectively by the user from 3D visualizations of the data. Thus, its choice would be extremely difficult to automate, and currently needs to be determined by the user. However, as will be illustrated in our experiments, γ needs to be roughly tuned only, and the same value can generally be used for all datasets of a same type. In the simplest cases, an appropriate value for γ may be found easily which provides a tradeoff between a suitable amount of denoising and a global enough interpolation, e.g. $\gamma = 1.5$ was obtained empirically to model brain ventricles from five different datasets acquired

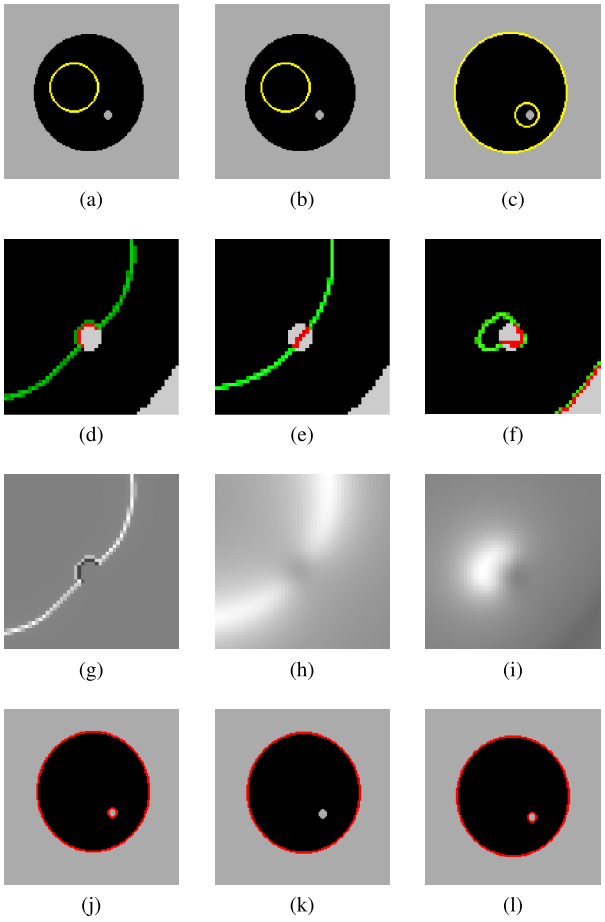

Fig. 8. Influence of γ and the initial position of the contour on the segmentation of a small hole illustrated via three examples. (a), (b), and (c) initial contour, (d), (e) and (f) zoom on front speed S at an intermediate state – green: positive; red: negative, (g), (h) and (i) zoom on $\frac{\partial \phi}{\partial t}$ at an intermediate state – light grey: positive; dark grey: negative, (j), (k) and (l) final segmentation. Left column: a sharp RBF ($\gamma = 0.1$) segments the small hole in the middle of the object. Middle column: a flat RBF ($\gamma = 2$, same initialization) treats the hole as noise and fails to delineate it, and Right column: initialization close enough to the edges of the small hole allows the speeds produced by the hole's edges to prevail against speeds produced by the object's points, resulting in the hole being segmented even with a flat RBF ($\gamma = 2$).

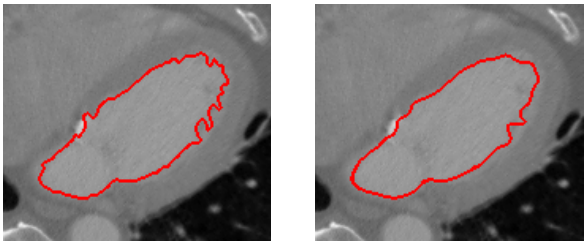

Fig. 9. Influence of the RBF flatness (via parameter γ) on the smoothing of surface details – the cavities of the left ventricle and atrium are segmented as one object. (a) a sharp RBF ($\gamma = 2$) captures the irregularities of the edge of the modeled object, while (b) a flat RBF ($\gamma = 6$) smoothes the segmentation and yields a less detailed contour. This image is a central long-axis slice of a CT-scan of the heart.

by two different MRI scanners at 1T and 1.5T. In other cases, we recommend two approaches, detailed below, which seek to provide a high degree of smoothing in the gaps

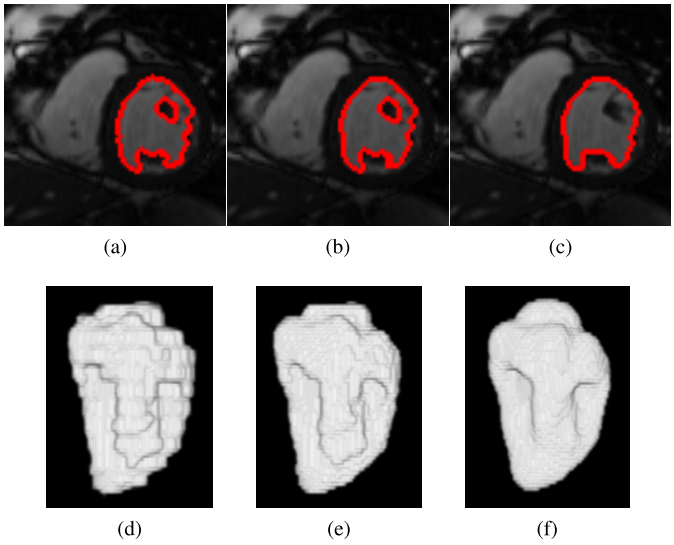

Fig. 10. Small object or hole preserving approach – the modeled object is the LV cavity of the heart of Fig. 2 with one central short-axis slice in the left column and 3D view in the right column. Processing of a dataset in two stages: (a) and (d) initial segmented and locally interpolated surface using a sharp RBF ($\gamma = 0.7$), and (b) and (e) final segmented and globally interpolated surface using a flat RBF ($\gamma = 3$), to be compared with (c) and (f) processing of the dataset using a flat RBF ($\gamma = 3$) only. The small object or hole preserving approach (two first columns) segments the papillary muscle inside the LV cavity, while the simplest approach with a single flat RBF ($\gamma = 3$), in the right column, treats it as noise.

while preserving the segmentation quality. The first approach preserves small objects and small holes which otherwise would be treated as noise when using a large value for γ. The second preserves also small details on the surface, e.g. small bumps, but at the expense of a slightly less satisfactory interpolation.

To obtain smooth and globally interpolated shapes while preserving small objects or holes – when a level set contour is initialized close to the edges of small objects or small holes, the speeds produced by these edges prevail against speeds produced by non-edge points, as illustrated in Figs. 8f and 8i, and the contour is therefore able to segment these small objects or holes (Fig. 8l). We generate such initializations automatically by first processing the data with a low value for γ, which yields the required level of denoising while preserving small objects and holes (Figs. 10a and 10d). This places the contour at the required initial position, i.e. on the edges of the modeled object and of the smaller objects and holes, e.g., as in Fig. 10a, with the left ventricle (LV) cavity and the papillary muscle that is inside it. Then, starting from the contour or surface just obtained, we process the data with a high value of γ in order to refine the interpolation (Figs. 10b and 10e). Note that in Fig. 10b, the papillary muscle is still found, while it was not when the same value for γ was used, without preliminary segmentation, in Fig. 10c, and hence it was erroneously included in the LV cavity.

To preserve fine details on the surface at the expense of slightly less smooth and globally interpolated shapes – to prevent fine surface details being lost when a particularly high degree of smoothing in the gaps is required, we suggest the application of several RBFs, each having a different flattening

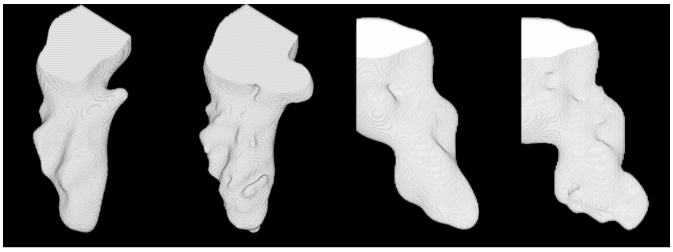

Fig. 11. Segmentation and interpolation of large gaps and small angle between the images and the modeled object's surface – the modeled objects are the LV cavity of a heart from sets of parallel long-axis CT-scan slices. Gap size: (a) and (b) 15 pixels, (c) and (d) 20 pixels. Reconstruction by, left column: the small object or hole preserving approach ($\gamma = 3$ then (a) $\gamma = 9$, and (c) $\gamma = 12$), and right column: the surface-detail preserving approach (γ ranging from 3 to (b) 12, and (d) 15).

strength, in the same interpolation, i.e.

$$\phi(\mathbf{x}) = \sum_{l=1}^{L} \sum_{\mathbf{x}_i \in \Omega_l} \alpha_i^l \psi^l(\mathbf{x} - \mathbf{x}_i), \tag{13}$$

where ψ^l is the RBF of flatness level l and Ω_l is the domain of the control points associated with RBFs of flatness level l, such that $\bigcup_l \Omega_l = \Omega$ and $\Omega_{l_a} \cap \Omega_{l_b} = \emptyset$ when $l_a \neq l_b$. These domains should be chosen such that sharp RBFs are used at data points, e.g. in the planes of 2D images, while, in the gaps, RBFs' flatness increases progressively away from data points. α_i^l is simply the previous coefficient α_i associated with the RBF ψ^l centered at point $\mathbf{x}_i$. We use this notation to establish that $\mathbf{x}_i$ and α_i are now associated with an RBF of a given flatness level l. Eqs. (7) and (9) then become:

$$\frac{d\alpha_i^{l_i}}{dt} = \left((\delta_\epsilon(\phi) \cdot S) \star \psi^{l_i} \right)(\mathbf{x}_i), \tag{14}$$

with l_i such that $\mathbf{x}_i \in \Omega_{l_i}$, and

$$\begin{aligned} \frac{\partial \phi}{\partial t}(\mathbf{x}) &= \sum_{l=1}^{L} \sum_{\mathbf{x}_i \in \Omega_l} \frac{d\alpha_i^l}{dt} \psi^l(\mathbf{x} - \mathbf{x}_i) \\ &= \sum_{l=1}^{L} \sum_{\mathbf{x}_i \in \Omega_l} H_l(\mathbf{x}_i) \frac{d\alpha_i^l}{dt} \psi^l(\mathbf{x} - \mathbf{x}_i) \\ &= \sum_{l=1}^{L} \left(\left(H_l \cdot \frac{d\alpha^l}{dt} \right) \star \psi^{l_i} \right)(\mathbf{x}) \end{aligned} \tag{15}$$

with H_l being equal to 1 inside Ω_l, and 0 elsewhere. When using this surface-detail preserving approach, the initial α can not be computed as easily as in (10). However, it is easy to update ϕ directly using (14) and (15), without keeping track of α. Although the resulting shapes will have a much higher degree of smoothing in the gaps than when using a sharp RBF only, they will be globally interpolated in the middle of the gaps only, where γ is the highest, while the interpolation will be more local close to the data, as illustrated in Fig. 11d. Hence, the resulting interpolation is less satisfactory than with the small object or hole preserving approach.

Fig. 11 illustrates the ability of the two proposed approaches to process the extreme case of sets of parallel slices having

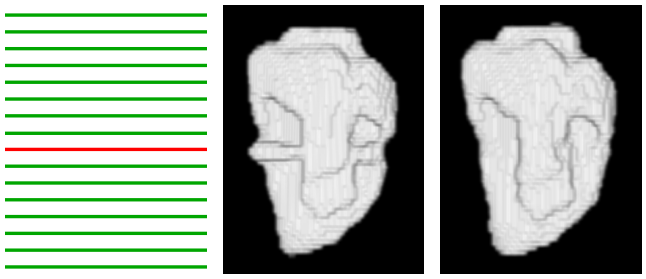

(a) (b) (c)

Fig. 12. Recovering after a local failure in the segmentation – the modeled object is the LV cavity of a heart of Figs. 2 and 10. (a) dataset viewed from the side – green: true slices; red: erroneous slice, (b) reconstructed object (same view), and (c) original object.

particularly large gaps and a small angle with the modeled object's surface, hence requiring an exceptionally high degree of smoothing in the gaps. In Fig. 11a, the small object or hole preserving approach smoothes out the fine details on the surface of the modeled object, while in Fig. 11b the surface-detail preserving approach is better at reconstructing them. On the other hand, in the dataset of Fig. 11c only four slices were available and thus did not contain much information on the fine details of the object's surface. Therefore, the small object or hole preserving approach could be used, with a very flat RBF ($\gamma = 12$) at the second step, without losing a significant amount of information. This approach yielded a better interpolation than the surface-detail preserving approach (Fig. 11d).

C. Choice of Segmentation Algorithm

The proposed framework may be used with any segmentation algorithm since no constraint is set on the computation of S. This choice of algorithm is dictated by the type of data and the imaging modality. Medical slices may be segmented using a number of image segmentation algorithms, such as edge based and region based methods, e.g. [21], [22] and [23], as well as methods that use prior knowledge, e.g. [24], [25].

When a dataset contains images with different gains and contrasts, or coming from different modalities, $S(\mathbf{x})$ can be computed independently for each image, using different segmentation algorithms. Then, the values in $S(\mathbf{x})$ for all images can be jointly used to compute $\frac{d\alpha_i}{dt}$ in (7). An example of the use of this approach may be found in Experiment 3, where T1 and T2 weighted MRIs of a brain are processed simultaneously.

The choice of the segmentation algorithm is as crucial to the success of interpolation as it is for segmentation, since interpolation relies on the position of the level set implicit surface in the image through the term S in (7). Thus, the outcome of the interpolation stage and of the integrated framework relies heavily on the choice of a robust segmentation algorithm by the user. It may be noted that the integration of interpolation and segmentation allows using 3D or 4D segmentations rather than the less robust independent 2D segmentations. Nevertheless, segmentation remains a difficult task, and is central to the success of the integrated framework. We note however that the failure of the segmentation in a part of the volume *has only an impact on the interpolation in the neighborhood* of this part, while in the rest of the volume where the segmentation is satisfactory, the framework recovers quickly. This is illustrated in Fig. 12 where a set of slices has been segmented and interpolated after replacing a central horizontal slice by a wrong one to simulate a failure in its segmentation.

D. Choice of Time Step

The proposed framework is very robust to numerical instabilities thanks to the use of an ODE rather than the traditional level set's PDE, and no numerical instabilities would arise, even for very large time steps, and without using any renormalization. We can therefore choose the time step depending on the maximum number of pixels the implicit surface is allowed to travel through in one iteration. Indeed, at a given position, a level set interface moves by distance D, defined as

$$D(\mathbf{x}) = \frac{dt \cdot \frac{\partial \phi}{\partial t}(\mathbf{x})}{\|\nabla \phi(\mathbf{x})\|}. \tag{16}$$

If D is too large, the implicit surface may pass through small objects without seeing them. In our experiments, we usually set $D = 1$ at the beginning of the process, which was then reduced automatically and progressively when oscillations of the implicit surface around the edges of the segmented object were detected. Lower D values allow the implicit surface to stabilize on object edges and to segment them more accurately.

IV. Experiments and Validation

To evaluate our proposed RBF based level set method, the interpolation aspect of the framework is assessed and compared with existing methods using both artificial and real medical data. In particular, we concentrate on the ability of the proposed method to handle sparse data by interpolating an object's shape in gaps, as this is the main contribution of this work. We do not evaluate the quality of the segmentation since this quality largely depends on the chosen segmentation algorithm, and this choice is highly dependent on the nature of the data. As the proposed framework is general enough to be used on data from a variety of different modalities (e.g. MRIs and CT-scans as shown in our experiments), and can accept any level set segmentation algorithm for the computation of the speed S of the level set implicit surface in (7), as outlined in Subsection III-C, we do not recommend any specific segmentation algorithm and let the user choose the most appropriate method depending on their data. Therefore, the choice of an optimum segmentation algorithm is beyond the scope of this work, and we only evaluate the quality of the interpolation.

Amongst several segmentation methods we experimented with, such as [21], [23] and [22], we selected the piecewise constant (PC) model used in [23] as the most suitable segmentation approach for validation of our framework on artificial data made of piecewise constant regions. For validation on real data, a method similar to [26] was used, which is based on a piecewise model and on a Parzen window method

to estimate the distributions of image intensity inside and outside of the object. However, we did not use the contour smoothing and implicit function regularization of [26] because our proposed method tends to produce smooth contours and does not require regularization, as mentioned in Section III-D. Also, we optimized the width of the Parzen window using log-likelihood maximization, while [26] set it equal to the standard deviations in the two regions. This was simple and generic enough to allow processing various types of images, e.g. both MRIs and CT-scans, with various amounts of noise.

We compare against the traditional sequential approach which performs interpolation and segmentation separately in turn, and our previous integrated method [8] which diffuses a level set segmenting surface in the gaps, but implemented with an SDF, rather than our original use of MCF. As explained in Section I, two sequential approaches exist: A) 2D segmentations of the image slices, followed by interpolation of the surface from the obtained binary masks, and B) interpolation of the images, followed by 3D segmentation of the obtained 3D volume. The type B sequential approach is more popular in automatic sequential processing of medical images, due to the more robust results produced by 3D segmentations over independent 2D segmentations of the slices. In addition, a fair comparison of the interpolation stage of the sequential and integrated approaches requires using the same segmentation stage in both methods. The proposed integrated framework can only use 3D and 4D segmentation by design, while sequential approach A is restricted to 2D independent segmentations. Due to these reasons, we compare against version B of the sequential approach, and we implement it with the state-of-the-art image interpolation method of Cordero *et al.* [7], and the same level set segmentation method as in our proposed integrated framework.

The image interpolation method in [7] uses non-rigid registration to find matching points in two images, and then interpolates the intensity between these points. The registration is obtained by optimizing, in multiresolution, the maximum a posteriori of a Markov Random Field (MRF), defined as the field of quantized deformations at control points. The MRF integrates intensity and gradient magnitude matching constraints between two corresponding areas of the registered images, a smoothness prior which penalizes large gradients of the transformation, and a topology preservation prior. Both forward and backward registrations are computed, in order to ensure that the whole procedure is unbiaised, and both results are used to compute the interpolation. Indeed, for each gap point, two sets (one for each forward and backward transformation) of two corresponding points are selected, so that the transformation path which links them passes through the gap point. The intensity at the gap point is then obtained as the linear combination of the intensities at the locations of the four matching points, weighted by terms that reflect the confidence in the selection of the image points.

We evaluate especially the ability of the methods to interpolate data having various spatial configurations, i.e. slices having arbitrary positions and orientations. This allows exploiting all available information without having to discard, for example, slices of unsupported orientations. Also, particular attention is paid to the quality of the interpolation in cases where a high amount of smoothing is required in the gaps. This is the case of gaps being comparatively too wide with respect to the size of the object and/or of slice orientations which do not match the principal orientation of the imaged object. Finally, we evaluate the capability of interpolating a variety of shapes, which can allow reconstructing either various objects or objects having a high degree of shape variability, e.g. some pathological organs such as brain ventricles of a patient suffering from hydrocephalus (see Fig. 17). Note that methods which integrate segmentation and interpolation in a model based framework are not general enough to allow processing shapes that have high variability, and we know of no other authors' work that *integrates* the whole process of segmentation and interpolation. Here, we compare against a typical sequential approach that we represent as one which consists of interpolation of the images first, e.g. using [7], and then segmentation in a second step (hereafter referred to as SM+[7] for "Sequential Method implemented with [7]"), and against our previous integrated segmentation and interpolation method [8], which was described in Section II.

A. Data and Method

We applied our proposed method to a wide range of real data of different modalities: (i) In Experiment 2 we used a collection of quasi-isotropic data from which a number of slices were removed in order to create gaps. These sets were one CT-scan of the acetabulum bone made up of parallel slices with pixels size of 0.49mm and slices spaced at 0.5mm, a CT-scan of the whole body with pixel size of 0.78mm and gaps of 0.8mm between the slices, from which we segment the LV cavity and kidney, and two MRI scans of a spherical phantom, with pixel sizes of 0.87mm and 1.44mm, and slice spacings of 0.9mm and 1.5mm respectively. (ii) In Experiment 3 we used five sparse MRIs of neonatal brains made of two stacks of T1 and T2 images, where each stack was composed of parallel images with a pixel size of 0.78mm and slices spaced at 3.3mm. We also used two sparse CT-scans of bones whose parallel slices had pixels of size 0.35mm and slices that were spaced at 0.5mm. (iii) In Experiments 4, 5 and 6, we used 16 sparse cardiac MRIs made of a stack of short-axis (SA) images and 0 to 12 long-axis (LA) images. The pixels ranged in size from 1.77mm to 2.08mm, while there were 5mm to 20mm gaps between the parallel SA images, and 25 time frames.

The interdependency of the segmentation and interpolation in the proposed framework makes it difficult to evaluate the interpolation stage alone and especially to assess its quality separately from the accuracy of the segmentation on real data. However, we note that as long as the segmentation is correct, the result of the interpolation stage depends only on the spatial configuration of the images and on the shape of the reconstructed object. For this reason, we evaluated the interpolation stage, in Experiment 1, and established its limits of performance in the case of a perfect segmentation, by segmenting and interpolating artificial data that is very simply and reliably segmented, e.g. images made of piecewise constant regions. We used the PC segmentation model used

in [21], which is intensity based, and we produced 3D datasets in which images are made of two regions of constant intensity, inside and outside of the object to model, respectively. One radial and one axial set were made by extracting radial and axial slices respectively from a volume containing a geometrical object made of a cylinder and a hemisphere. Then, 32 more sets were made from the cardiac MRI datasets of Experiment 4. First, two models of the LV cavity were generated by segmenting, using the proposed method, two real cardiac MRI datasets containing almost no gaps, hence requiring only a minimal amount of inter-slice interpolation. Then, these two models were used to generate 16 new datasets each, by using the positions and orientations of the 2D slices of our 16 cardiac MRI datasets and by replacing their intensities by those of the models. Therefore, even if the slices of these 32 sets are black and white images, they contain the shape of real hearts. Their positions and orientations are also realistic since they are taken from real cardiac MRIs. See Fig. 13 for some examples.

For comparative evaluation against [8] and SM+[7], we measured accuracy on the 3D shape of the segmented and interpolated objects (again because the proposed method and our previous method [8] cannot separate segmentation from interpolation). For better parity, the segmentation algorithm of our framework was used to segment the interpolated volume produced by Cordero *et al.* [7]. We used the Jaccard coefficient as an accuracy measure of the combined segmentation and interpolation. In Experiment 1, the groundtruth was the models used to produce the artificial data. In Experiment 2, in order to assess the accuracy of the interpolation on real data, we compared the reconstruction of the objects from the datasets with gaps against the reconstructions from the original full volume datasets. The same segmentation method was used in both cases to ensure that the segmentation results are similar and that the comparisons only evaluate the accuracy of the interpolation. In Experiments 3, 4, 5 and 6, we perform subjective comparisons on the non-isotropic real datasets. The values of γ used in the experiments are given in Table III.

TABLE III
VALUES OF PARAMETER γ USED IN THE EXPERIMENTS

Experiment	Modeled object	Spatial configuration	γ
Experiment 1	Geometric model	All	3
	LV models A and B	All	0.5 then 4
Experiment 2	Phantom Low Res.	All	20
	Phantom High Res.	All	20
	Acetabulum	All	9
	Kidney	5 pix. spacing	3 to 7
		10 pix. spacing	3 to 9
		15 pix. spacing	3 then 9
		20 pix. spacing	3 then 12
	LV cavity	5 pix. spacing	3 to 7
		10 pix. spacing	3 to 9
		15 pix. spacing	3 to 12
		20 pix. spacing	3 then 12
Experiment 3	Vertebrae	All	1.5
	Brain ventricles	All	1.5
Experiment 4	LV cavity	All	0.7 then 3
Experiment 5	4D LV cavity	All	2 ($\beta = 4$)
Experiment 6	LV + RV	All	0.7 then 3

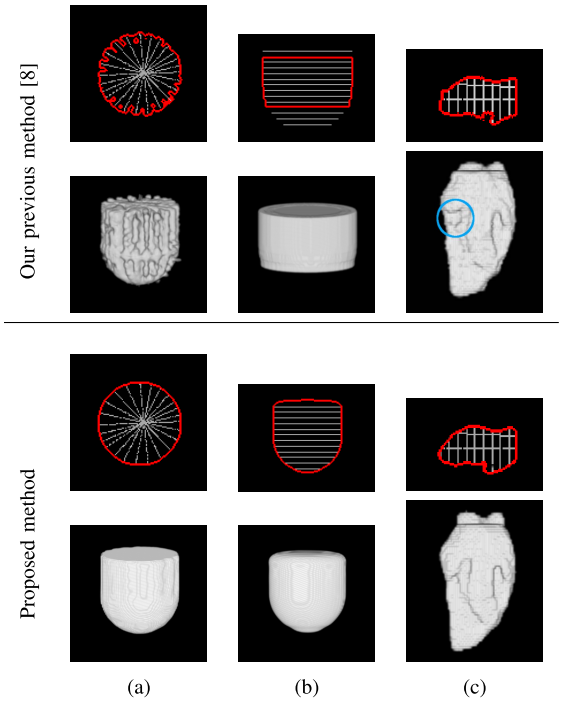

Fig. 13. Segmentation and interpolation of 3D artificial data. (a) radial dataset – top: central horizontal slice, bottom: 3D view, (b) axial dataset – top: central vertical slice, bottom: 3D view, (c) LV cavity model – top: central vertical slice, bottom: 3D view. The blue circle highlights a failure of [8] to recover the topology of part of the LV cavity.

B. Experiment 1: Quantitative Analysis on Artificial 3D Data

For the proposed method and [8], we segmented and interpolated the artificial 3D datasets starting from a small spherical surface located at the centre of the volume, using the PC model used in [23] to drive the segmentation.

Fig. 13 presents the segmented and interpolated surfaces of the radial and axial datasets and of an example of a modified cardiac MRI. Table IV provides the Jaccard coefficients indicating the proportions of voxels correctly assigned to the objects. For the two LV models, the coefficients are averaged over the 16 datasets made from each model. In the case of the proposed method, the Jaccard coefficients were obtained using the small object or hole preserving approach presented in Subsection III-B, with γ being equal to 0.5 and 4 in the first and second steps respectively, for all LV datasets. Even when no attempt was made to tune γ to match the individual spatial configurations of the datasets, the proposed method yielded better results than our previous method [8], with average accuracy measures of 0.943 against 0.906 for our previous method. We obtained a more global interpolation than our previous method [8] on both the geometrical and LV datasets, and managed to reach all the slices of the axial set, obtaining an accuracy of 0.967 (Fig. 13a), while [8]'s diffusion method optimized the contour only locally, failing to recover

TABLE IV
JACCARD COEFFICIENTS ON ARTIFICIAL 3D DATA

Dataset	Our previous method [8]	Proposed method
Radial	0.775	0.984
Axial	0.524	0.967
LV A	0.931 ± 0.022	0.951 ± 0.018
LV B	0.913 ± 0.021	0.931 ± 0.020
All LV	$\mathbf{0.922 \pm 0.023}$	$\mathbf{0.941 \pm 0.021}$
All datasets	$\mathbf{0.906 \pm 0.075}$	$\mathbf{0.943 \pm 0.022}$

the topology of some parts of the LV cavity properly (blue circle in Fig. 13c), and could not reach images located further away than its narrowband width in the axial data set, resulting in an accuracy of only 0.524 (Fig. 13b).

For the proposed method, we identified two difficult types of surface areas whose reconstructions most contributed in reducing the Jaccard accuracy measure. The first type is the centres of gaps where the object's surface has a sharp angle with the image planes, as was illustrated in Fig. 7, and which require a flat RBF in order to compensate for the increased distance between the data points and obtain a global and smooth interpolation. For example, this is the case for gaps near the apex of the LV models and at the bottom of the hemisphere of the axial dataset. Note that very large gaps would have the same effect, as will be illustrated in Experiment 2. As discussed in Section III-B, a tradeoff must be found in this case, between a flat RBF which takes into account the global geometry of the object, and a sharper RBF which better preserves the surface's details. In this experiment, the small object or hole preserving approach achieved this tradeoff satisfactorily for the LV models, thus minimizing the decrease of the Jaccard measure, and a high value for γ could be used with the axial dataset. The second type is the disappearance of a structure between two slices where the proposed method has no other choice than to guess the shape of the missing extremity, and it does so by attempting to preserve the smoothness of the global shape, and therefore produces plausible, yet possibly wrong, reconstructions. This is particularly true for large gaps where the amount of missing information is high. Overall, we found that the reductions of the Jaccard accuracy measure were due to the accumulation of several small deviations from the groundtruth surface, scattered over all the surface and with a slightly higher concentration in these two types of difficult areas. Note that the radial dataset does not suffer from any of these difficulties, and achieved the best Jaccard coefficient.

C. Experiment 2: Quantitative Analysis on Real 3D Data

We also evaluated our framework on real quasi-isotropic datasets from which a number of slices were removed in order to create gaps. We made 8 datasets from each original set, 4 being made of parallel slices spaced by 5, 10, 15 and 20 pixels respectively, and 4 with the same horizontal slices plus one additional vertical slice in the center of the volume. In order to demonstrate the versatility of the proposed method, we performed these tests on different objects having various shapes and sizes, and imaged using different modalities. The same segmentation algorithm, based on pixel intensity classification using a Parzen window estimator, similar to [26], was used for all compared methods.

TABLE V
JACCARD COEFFICIENTS ON REAL 3D DATA

Dataset	Slice spacing (pixels)	SM+[7]	Proposed method	
			Horz.	Horz.+Vert.
Phantom Low res.	5	0.970	0.989	0.990
	10	0.940	0.991	0.989
	15	0.923	0.978	0.987
	20	0.863	0.962	0.977
Phantom High res.	5	0.993	0.992	0.990
	10	0.980	0.991	0.993
	15	0.949	0.990	0.992
	20	0.917	0.985	0.990
Kidney	5	0.955	0.953	0.955
	10	0.895	0.887	0.900
	15	0.840	0.845	0.869
	20	0.771	0.802	0.846
LV	5	0.966	0.962	0.962
	10	0.919	0.916	0.922
	15	0.872	0.888	0.899
	20	0.750	0.835	0.864
Acetabulum	5	0.972	0.981	0.975
	10	0.948	0.968	0.968
	15	0.892	0.929	0.935
	20	0.881	0.903	0.915
Average of all datasets	**5**	$\mathbf{0.971 \pm 0.014}$	$\mathbf{0.975 \pm 0.017}$	$\mathbf{0.974 \pm 0.016}$
	10	$\mathbf{0.936 \pm 0.032}$	$\mathbf{0.951 \pm 0.047}$	$\mathbf{0.954 \pm 0.041}$
	15	$\mathbf{0.895 \pm 0.043}$	$\mathbf{0.926 \pm 0.061}$	$\mathbf{0.936 \pm 0.054}$
	20	$\mathbf{0.836 \pm 0.072}$	$\mathbf{0.897 \pm 0.079}$	$\mathbf{0.918 \pm 0.065}$
	All	$\mathbf{0.910 \pm 0.066}$	$\mathbf{0.937 \pm 0.059}$	$\mathbf{0.946 \pm 0.049}$

We compare the reconstructions from the datasets that have gaps against the reconstructions from the original full volume datasets in order to assess the quality of the interpolation. As already illustrated in Experiment 1, our previous method [8] does not scale correctly to large datasets and fails to propagate the contour in large gaps, so we only compare against SM+[7] in this experiment.

Table V provides the Jaccard coefficients for all of the datasets for SM+[7] and the proposed method when the datasets were made up of parallel horizontal slices only (in the third and fourth columns of the table), and parallel horizontal slices plus one vertical slice (in the last column).

The LV and Kidney datasets required a particularly large amount of smoothing in the gaps because of the combined effects of very large gaps and the unmatched orientations of (all) parallel slices against the orientation of the organs, hence causing those slices to have a small angle with the object surface. Therefore, for the proposed method, they were reconstructed using the surface-detail preserving approach for the smaller gap sizes (5–15 and 5–10 pixels for the LV and Kidney datasets respectively) and the small object or hole preserving

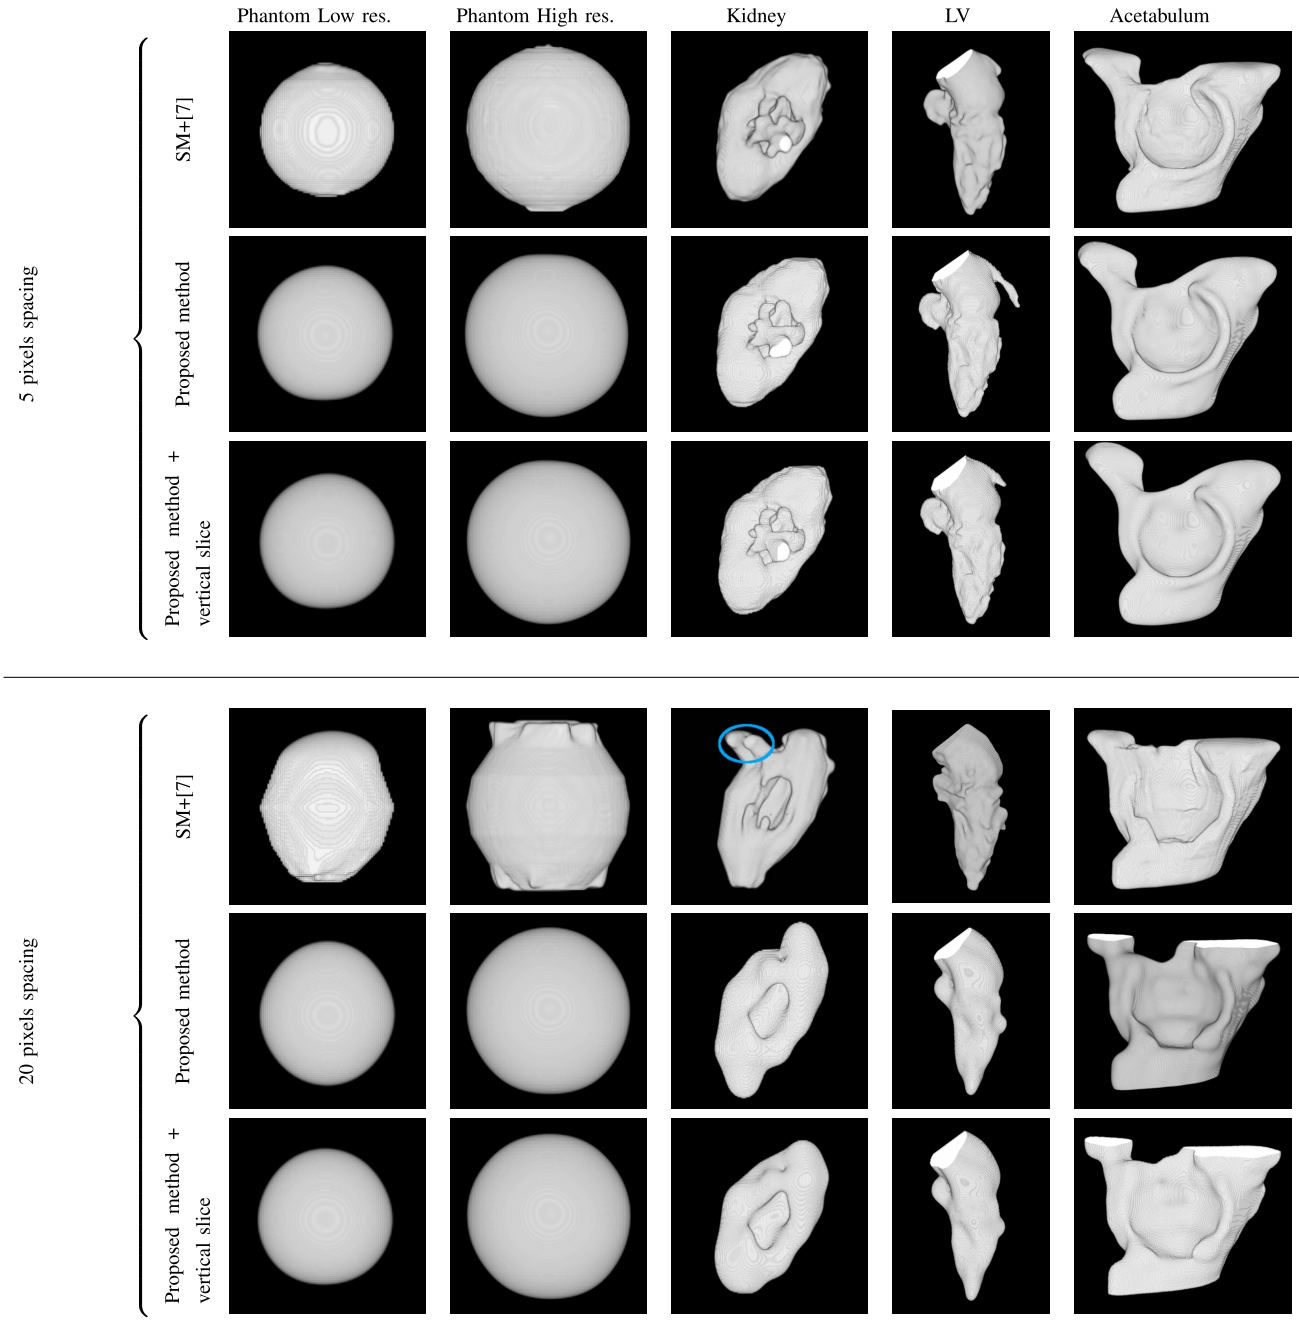

Fig. 14. Segmentation and interpolation of various shapes from 3D real data at two different slice spacings (5 and 20 pixels respectively), and from a variety of modalities. The blue circle highlights how [7] struggles to reconstruct extremities of objects for large slice spacings, especially where a background of similar intensity can be confused with the object.

approach for the larger gaps. γ was tuned depending on the gap size, ranging from 3 to 12, as shown in Table III. The phantom and acetabulum datasets presented smoother surfaces, therefore a simpler scheme was used for them, with a single value of γ for all slice spacing. In these cases, γ was only roughly adapted to the size of the object and the level of detail of its shape, and was set to 9 for the acetabulum, and 20 for the two types of phantom. For the method of Cordero *et al.* [7], we used the three default parameters provided by the authors, as we found that modifying them did not significantly change the interpolation followed by segmentation results.

Even when γ was not finely tuned for all datasets, the proposed method performed similarly or better than SM+[7]. While the segmentations of the slices, in the planes of the images, were subjectively similarly good for the two compared methods for all slice spacings, the interpolation by the proposed method produced better results in the gaps between the images, especially for larger gaps, with average accuracy measures of 0.937 and 0.946 across all datasets and all spacings, for parallel horizontal slices and horizontal plus vertical slices respectively. Cordero *et al.* [7] only managed a rate of 0.910 for parallel horizontal slices (and can not process datasets that have vertical slices).

Cordero *et al.* [7] interpolates linearly between two neighboring slices, thus producing a local interpolation of the shape, while our proposed method can take into account more than

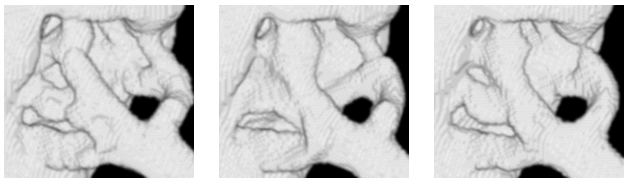

(a) (b) (c)

Fig. 15. Recovery of the topology of vessels from the Kidney dataset with 5 pixels spacing. (a) original dataset processed by (b) SM+[7], and (c) proposed method.

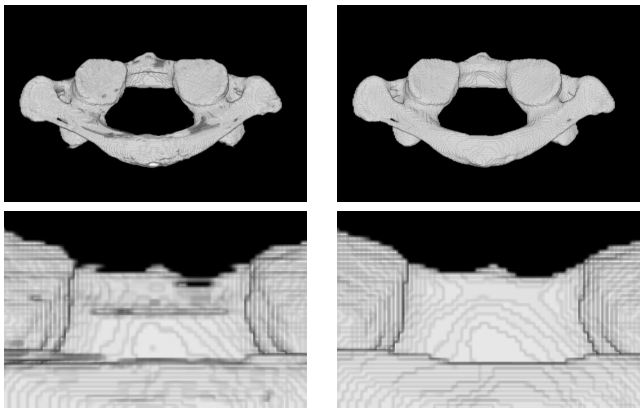

(a) (b)

Fig. 16. 3D segmentation and interpolation of the c1 vertebrae of Figs. 3a and 3b from a CT-scan. (a) our previous method [8], (b) proposed method. Top row: global view, bottom row: zoomed view.

two slices, especially for large values of γ, and therefore it provides a more global interpolation. This is illustrated in the 4th row of Fig. 14 where the results of Cordero *et al.* [7] for 20 pixel spacing are made of straight segments because of the linear and local properties of the interpolation and yield accuracy rates ranging from 0.750 to 0.917. On the contrary, in the 5th and last rows of Fig. 14, the proposed method gave a smooth and rather round shape, and accuracy rates of 0.802 to 0.985 and 0.846 to 0.990 respectively. Cordero *al.* [7] also struggled where the object disappears between two slices, such as at the extremities of the phantom scans. This resulted in protrusions at these positions – for example, their method confused the object with a background of similar intensity in the 20 pixel spacing Kidney dataset (see blue circle in the 4th row of Fig. 14) which gave an accuracy rate of only 0.771.

The addition of a single vertical slice increased the robustness of the proposed method in large gaps (right column of Table V and 3rd and 6th rows of Fig. 14), with average accuracy rates across all five datasets of 0.897 to 0.918 for 20 pixel spacings, even if the position and orientation of the additional slice were not chosen to match to the position and orientation of the imaged object but instead were selected arbitrarily. Recall, Cordero *et al.*'s approach [7] is inherently unable to use LA images.

Fig. 15 illustrates another example of shape detail extraction where, for the Kidney dataset with 5 pixels spacing, we can see that the topology of the vessels recovered by the proposed method is much closer to the original volume than Cordero *et al.* [7]'s result, which fused some of the vessels together.

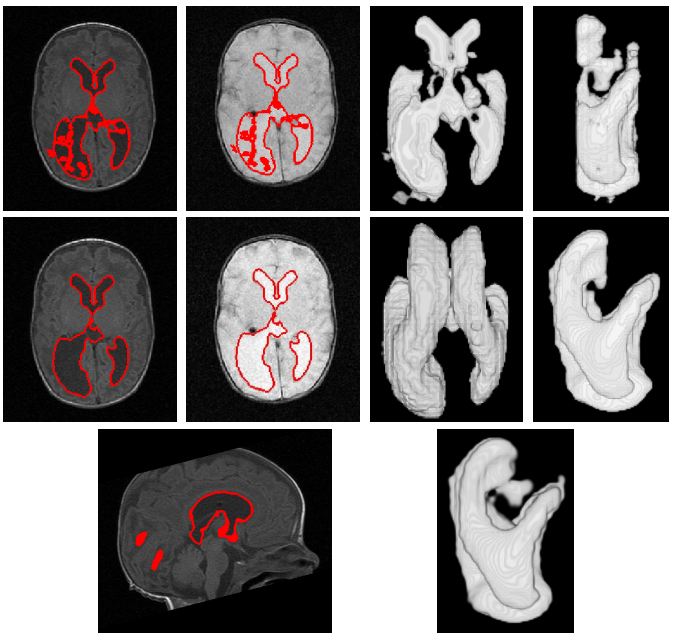

Fig. 17. 3D segmentation and interpolation of the ventricles of the neonatal brain of Fig. 3c from combined T1 and T2 MRIs. Top row: our previous method [8], middle row: proposed method. From left to right: T1 central slice, T2 central slice, axial 3D view, sagittal 3D view. Bottom row: Proposed method which additional sagittal set of slices. Left: central sagittal slice, right: sagittal 3D view.

D. Experiment 3: Qualitative Analysis on Real CT-Scans and MRIs

In this experiment, vertebrae were segmented from CT-scans (Fig. 16), and the ventricles of the brain were segmented from neonatal brain MRIs (Fig. 17), using the Parzen estimator based segmentation algorithm. The datasets were visually inspected for any visible misalignment as a pre-processing step and also considered not to require much smoothing (unlike the data used in Experiment 2), because the gaps were smaller and the slices were acquired roughly perpendicular to the principal axes of the modeled objects. Therefore, the simplest scheme with a single value of $\gamma = 1.5$ was used for all the datasets of this experiment.

We did not compare against SM+[7] in this experiment, because their method requires images spaced by at least twice their pixel size, which was not the case for the vertebra CT-scans, and it can not handle images from different modalities simultaneously, such as the T1 and T2 images of the brain MRIs. On the contrary, the proposed method could jointly process the T1 and T2 images, by applying different segmentation algorithms to the computation of S for each modality, as suggested in Section III-C. Note that some methods perform mutimodal segmentation on brain MRIs with small gaps by "stacking" the image slices on top of each other, ignoring the gaps. The drawback of such methods is that they would discard any slice with a different orientation, such as sagittal and coronal slices, if available. By preserving the original spacing between the axial slices, our proposed framework can make use of these different slices. In Fig. 17, this allowed using an additional set of sagittal slices, which helped to segment a feature which was missed when using the axial slices only.

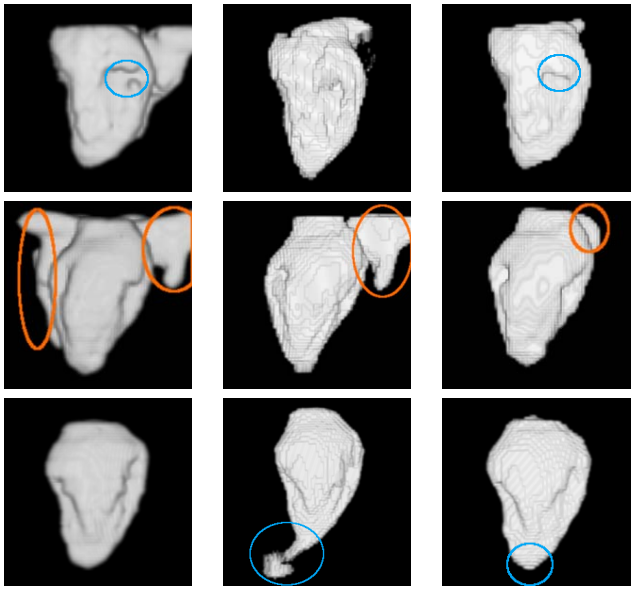

(a) (b) (c)

Fig. 18. 3D segmentation and interpolation of the LV cavity from three different cardiac MRIs. (a) our previous method [8], (b) SM+[7], (c) proposed method. The extraneous regions protruding from some LV cavities for all methods in some images, e.g. orange circles in the second row above, are parts of the right ventricle which have been wrongly segmented by the chosen segmentation algorithm. The blue circles highlight situations where the proposed method recovered a better topology than [8] and [7].

It can be seen in Fig. 16 that our previous method [8] did not interpolate the bone's shape in the gaps satisfactorily, because of its local approach. However, the proposed method produced a much more satisfactory global interpolation. Similarly, in Fig. 17, [8] failed to extend the contour to all the parts of the brain ventricles due to the presence of large gaps in the data. In addition, the high degree of noise in the brain MRI datasets biased its segmentation. The proposed method did not suffer from these issues and the segmenting contour reached all the parts of the brain's ventricles (Fig.17). The segmentation was not impaired by the high level of noise, due to the inherent robustness of the method to noise, as described earlier in Section III-B.

E. Experiment 4: Qualitative Analysis on Real Cardiac MRIs

We modeled the cavity of the LV (Fig. 18) from cardiac MRIs, after visual inspection for no visible misalignment. The clear intensity separation between blood-pool and muscle allowed using the simple PC segmentation model. The datasets required about the same degree of smoothing as those of Experiment 3, but the simplest scheme with a single value for γ failed to delineate the papillary muscles. Therefore, we used the small object or hole preserving approach presented in Subsection III-B, with $\gamma = 0.7$ and $\gamma = 3$ for the first denoising stage and the second interpolation refinement stage respectively, for all the datasets. The RBF of the latter stage was sharp enough to preserve the details on the surface of the modeled objects. The image interpolation method of Cordero *et al.* [7] was ran and tuned by the authors of [7]. The cardiac MRIs are made of 25 time frames, but only one time frame was used, as methods [7], [8] were designed for 3D segmentation and interpolation only. In addition, Cordero *et al.* [7] used only SA slices from these sets and discarded all LA slices.

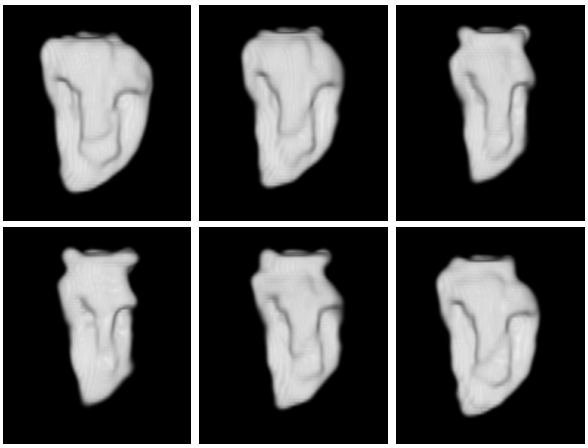

Fig. 19. 4D segmentation and interpolation of the LV cavity from a cardiac MRI: 6 different time frames.

When subjectively compared, we observe that the segmentations of the 2D images of the cardiac datasets were similarly accurate for all three methods, but the interpolation was better using the proposed method, as illustrated by the blue highlighting circles in Fig. 18. Notably, the apex of the LV is better segmented and interpolated by the proposed method than by SM+[7] thanks to the use of the LA slices.

F. Experiment 5: Real 4D Data

An example of a 4D segmentation and interpolation performed by the proposed method is shown in Fig. 19. As in Experiment 4, we used the PC segmentation model. This cannot be compared with [7], [8] since they do not handle 4D data as mentioned in the previous experiment. Cordero *et al.* [7] can perform interpolation of the volume in each time frame, but not across the time domain in an interframe fashion.

In Section III, FFTs were used to compute the convolutions of (7), (14) and (15). As an additional bonus, FFT's inherent interpretation of periodic data allowed us to exploit the cyclical nature of cardiac motion. Joint segmentation, interpolation and tracking through time from non-periodic signals may be done by padding the data with a number of blank frames.

G. Experiment 6: Multiple Regions

This experiment aims to demonstrate the possibility to extend the framework to several regions. Multi-region segmentation is a common feature of level set methods, which the proposed level set framework naturally inherits from. Multi-region level set schemes have been proposed in the past, e.g. by Vese and Chan [27], who used simultaneously a number N of level sets in order to segment 2^N piecewise constant or piecewise smooth regions. The regions are defined as all the possible intersections of the areas bounded by the different level sets, and, by construction, they cannot overlap or create vacuum. We used this region definition in order to segment simultaneously the cavities of the left and right ventricles (RV) of the heart and the myocardium from cardiac cine MRIs. In Fig. 20a, the two areas which are inside one level set contour

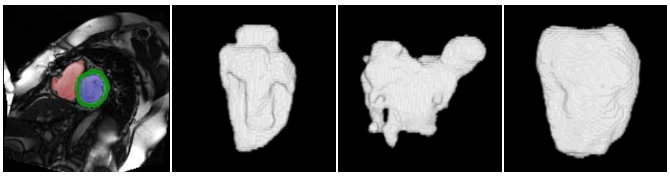

(a) (b) (c) (d)

Fig. 20. Segmentation of multiple regions using the proposed framework. (a) SA slice with the LV, RV and myocardium segmentations colored in blue, red, and green respectively. 3D reconstructions of (b) the LV, (c) the RV, (d) the myocardium.

and outside of the other (blue and red areas) are the LV and RV cavities (Figs. 20b and 20c respectively). The myocardium is the area where both level set functions are positive (green area in Fig. 20a, and Fig. 20d). In this example, the PC method was used for the segmentation. However, it was not robust enough to segment the myocardium, whose intensities are very similar to these of the background. Therefore, a constraint has been added to the thickness of the myocardium in order to prevent the contour from leaking to the background. We used a method similar to what was presented in [28] for edge based segmentation, but adapted to the PC segmentation algorithm. An initial segmentation of the two LV and RV cavities is computed in a first step, using two level sets which are updated simultaneously and do not need to intersect and interact as a result of the natural separation of the two cavities. Then, the segmentation of the myocardium is added to the process, driven by both a coupling term $S_{coupling}$ between the endo-cardium and the epi-cardium, and an image derived term S_{image}, limited to the area of expected location of the epi-cardium:

$$S = \lambda S_{coupling} + (1-\lambda) S_{image} \ . \tag{17}$$

λ is a weighting coefficient and is set to 0.5 in our test. The coupling term is computed as:

$$S_{coupling} = \begin{cases} -1 & \text{if } \phi_1 \leq -d - w \\ \sqrt[3]{\frac{d-|\phi_1|}{w}} & \text{if } -d-w < \phi_1 < -d+w \\ & \text{and } S_{PC} > 0 \\ S_{PC} & \text{if } -d-w < \phi_1 < -d+w \\ & \text{and } S_{PC} \leq 0 \\ S_{PC} & \text{if } \phi_1 \geq -d+w \end{cases} . \tag{18}$$

ϕ_1 is the value of the level set function associated with the LV cavity contour (or endo-cardium) at the end of the initial segmentation stage, and therefore it represents the (negative) distance to this contour. d and w are the nominal distance between the endo- and epi-cardium and the transition width respectively. In [28] they are derived from an a priori model, but for simplicity, they are set to 6 and 4 in our test. S_{PC} is the speed yielded by the PC segmentation model. The image derived term is defined as:

$$S_{image} = \begin{cases} \exp\left(-\frac{(x-d)^2}{2w}\right) S_{PC} & \text{if } \phi_1 \leq -d \text{ and } S_{PC} > 0 \\ S_{PC} & \text{otherwise.} \end{cases} \tag{19}$$

H. Timing

Our framework was implemented in C++, and the experiments were run on a 1.6Ghz Unix machine. It is to be noted that timing comparisons against SM+[7] mainly highlight differences in the interpolation time, since our implementation of the level set framework was used for the segmentation stage of SM+[7]. Therefore, as will be detailed next, the main differences arise from the additional interpolation step of SM+[7], and varying numbers of iterations.

In the cases where the simple scheme with a single value for γ could be used, the proposed method segmented and interpolated the data in about 5 to 15 minutes, depending on the size of the dataset. SM+[7] was moderately slower by approximately 5 to 20 minutes, since the running time of [7] had to be added to the segmentation time, and the segmentation itself (using the same segmentation framework as for the proposed method) was slightly longer due to the higher number of data points.

Conversely, when two values for γ were used, due to the dataset requiring more interpolation, the proposed method needed more iterations and time (30 minutes to 2 hours) to process the data. Indeed, the second, interpolation refining step, required 4 to 10 times more iterations than the segmentation step in order to fully propagate the global interpolation in the gaps and replace the existing local one. Note that these computation times are still reasonable for the intended applications, which do not require real time processing. SM+[7] performed in the exact same way as in the first case, with the same number of data points and iterations required, and was therefore faster than the proposed approach for this configuration.

V. Conclusion

We presented an integrated segmentation and interpolation framework to reconstruct 3D and 4D sparse datasets based on a novel RBF interpolated level set methodology. This new framework can handle sets of 2D slices having any spatial configuration, i.e. made of any number of slices having arbitrary positions and orientations. The integrated framework was validated on artificial and real data, including MRIs and CT scans. The method proved to be very robust, thanks to the interdependency of the segmentation and interpolation processes, and produced better results than our previous integrated method [8] and the sequential approach implemented with [7]. The proposed method is not affected by differences in gains and contrasts in the image slices and is able to combine different types of images in the same processing, e.g. T1 and T2 MRIs, provided they are properly aligned. For future work, we plan to investigate the potential for the application of the new framework to object reconstruction from sparse and noisy 3D point clouds such as those produced by the Microsoft Kinect sensor. We also wish to add registration to this framework in order to handle sets of misaligned slices, or possibly several point clouds simultaneously, for increased signal quality. The inclusion of registration will be mutually beneficial to all the elements of the framework, i.e. segmentation, interpolation, and registration.

ACKNOWLEDGMENT

We would like to thank the authors of [7] who allowed us to compare to their method by kindly both running their software on our data and making their code available for further experiments. The CT-scans datasets used in the experiments were from the Laboratory of Human Anatomy and Embryology, University of Brussels (ULB), Belgium.

REFERENCES

[1] G. Li, S. Lu, Z. Min, J. Li, X. Xu, and E. Song, "3D prostate boundary reconstruction from 2-D TRUS images," in *Proc. Bioinformatics Biomed. Eng., Int. Conf.*, 2007, pp. 940–943.

[2] A. Pihuit, O. Palombi, and M.-P. Cani, "Reconstruction implicite de surfaces 3-D à partir de règions 2-D dans des plans Parallèles," in *Proc. AFIG*, 2009, pp. 2–4.

[3] L. Liu, C. Bajaj, J. O. Deasy, D. A. Low, and T. Ju, "Surface reconstruction from non-parallel curve networks," *Comput. Graph. Forum*, vol. 27, no. 2, pp. 155–163, 2008.

[4] Z. Pan, X. Yin, and G. Wu, "Segmentation-based interpolation of 3-D medical images," in *Proc. ICCSA*, 2004, pp. 731–740.

[5] S. U. Rahman and S. Wesarg, "Combining short-axis and long-axis cardiac MR images by applying a super-resolution reconstruction algorithm," *Proc. SPIE*, vol. 7623, pp. 76230I-1–76230I-12, Mar. 2010.

[6] X. Yuan and X. Yuan, "Fusion of multi-planar images for improved three-dimensional object reconstruction," *Comput. Med. Imag. Graph.*, vol. 35, no. 5, pp. 373–382, 2011.

[7] L. Cordero-Grande, G. Vegas-Sanchez-Ferrero, P. Casaseca-de-la Higuera, and C. Alberola-Lopez, "A Markov random field approach for topology-preserving registration: Application to object-based tomographic image interpolation," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 2047–2061, Apr. 2012.

[8] A. Paiement, M. Mirmehdi, X. Xie, and M. Hamilton, "Simultaneous level set interpolation and segmentation of short and long-axis MRI," in *Proc. MIUA*, 2010, pp. 267–271.

[9] J. Woo, E. Murano, M. Stone, and J. Prince, "Reconstruction of high-resolution tongue volumes from MRI," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 12, pp. 3511–3524, Sep. 2012.

[10] H. van Assen, M. Danilouchkine, M. Dirksen, J. Reiber, and B. P. F. Lelieveldt, "A 3-D active shape model driven by fuzzy inference: Application to cardiac CT and MR," *IEEE Trans. Inf. Technol. Biomed.*, vol. 12, no. 5, pp. 595–605, Sep. 2008.

[11] S. Zambal, J. Hladuvka, and K. Bühler, "A fully automatic system for segmentation and analysis of the left and right ventricles of the heart using a bi-temporal two-component model," in *Proc. CARS*, 2007, pp. 93–94.

[12] J. M. Lötjönen, V. M. Järvinen, B. Cheong, E. Wu, S. Kivistö, J. R. Koikkalainen, *et al.*, "Evaluation of cardiac biventricular segmentation from multiaxis MRI data: A multicenter study," *J. Magn. Resonance Imag.*, vol. 28, no. 3, pp. 626–636, 2008.

[13] L. Cordero-Grande, G. Vegas-Sànchez-Ferrero, P. C. de-la Higuera, J. A. San-Romàn-Calvar, A. Revilla-Orodea, M. Martín-Fernández, *et al.*, "Unsupervised 4-D myocardium segmentation with Markov random field based deformable model," *Med. Image Anal.*, vol. 15, no. 3, pp. 283–301, 2011.

[14] D. L. Chopp and J. A. Sethian, "Flow under curvature: Singularity formation, minimal surfaces, and geodesics," *Experim. Math.*, vol. 2, no. 4, pp. 235–255, 1993.

[15] D. L. Chopp and J. A. Sethian, "Motion by intrinsic Laplacian of curvature," *Inter. Free Boundaries*, vol. 1, no. 1, pp. 107–123, 1999.

[16] R. Franke, "Scattered data interpolation: Tests of some method," *Math. Comput.*, vol. 38, no. 157, pp. 181–200, 1982.

[17] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, *et al.*, "Reconstruction and representation of 3-D objects with radial basis functions," in *Proc. Comput. Graph. Interact. Tech. SIGGRAPH*, 2001, pp. 67–76.

[18] S. Wang, K. Lim, B. Khoo, and M. Wang, "An extended level set method for shape and topology optimization," *J. Comput. Phys.*, vol. 221, no. 1, pp. 395–421, 2007.

[19] G. Slabaugh, Q. Dinh, and G. Unal, "A variational approach to the evolution of radial basis functions for image segmentation," in *Proc. IEEE Conf. CVPR*, Jun. 2007, pp. 1–8.

[20] O. Bernard, D. Friboulet, P. Thevenaz, and M. Unser, "Variational B-spline level-set: A linear filtering approach for fast deformable model evolution," *IEEE Trans. Image Process.*, vol. 18, no. 6, pp. 1179–1191, Jun. 2009.

[21] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *Int. J. Comput. Vis.*, vol. 22, no. 1, pp. 61–79, 1997.

[22] R. Yang, M. Mirmehdi, and X. Xie, "A charged active contour based on electrostatics," in *Proc. ACIVS*, 2006, pp. 173–184.

[23] T. Chan and L. Vese, "Active contours without edges," *IEEE Trans. Image Process.*, vol. 10, no. 2, pp. 266–277, Feb. 2001.

[24] M. Leventon, W. E. L. Grimson, and O. Faugeras, "Statistical shape influence in geodesic active contours," in *Proc. IEEE Conf. CVPR*, vol. 1. Jun. 2000, pp. 316–323.

[25] D. Cremers, "Nonlinear dynamical shape priors for level set segmentation," in *Proc. IEEE Conf. CVPR*, Jun. 2007, pp. 1–7.

[26] Y. Li, C. Li, and Y. Zhang, "Novel level set model for the image segmentation based on Parzen-Window," in *Proc. IEEE Conf. CSAE*, vol. 3. May 2012, pp. 217–221.

[27] L. A. Vese and T. F. Chan, "A multiphase level set framework for image segmentation using the Mumford and Shah model," *Int. J. Comput. Vis.*, vol. 50, no. 3, pp. 271–293, 2002.

[28] M. Lynch, O. Ghita, and P. Whelan, "Left-ventricle myocardium segmentation using a coupled level-set with a priori knowledge," *Comput. Med. Imag. Graph.*, vol. 30, no. 4, pp. 255–262, 2006.

Adeline Paiement received the Diplôme d'ingénieur degree from Ecole Nationale Supérieure de Physique de Strasbourg, France, and the M.Sc. degree in physics (with commendation) from Université Louis Pasteur, Strasbourg, in 2008. She is currently pursuing the Ph.D. degree with a focus on analysis of 3-D/4-D sparse data, including medical images with the Visual Information Laboratory, University of Bristol, Bristol, U.K. Her research interests include image/volume registration, segmentation and interpolation, level set methods, and medical imaging.

Majid Mirmehdi received the B.Sc. (Hons.) and Ph.D. degrees in computer science in from the City University, London, U.K., 1985 and 1991 respectively. He is currently a Professor of computer vision with the Department of Computer Science, University of Bristol, and the Graduate Dean and Head of the Graduate School of Engineering. His research interests include natural scene analysis and medical imaging, and he has published more than 150 refereed conference and journal publications in these and other areas. He is a fellow of the International Association for Pattern Recognition. He is Editor-in-Chief of the *IET Computer Vision Journal* and an Associate Editor of the *Pattern Analysis and Applications Journal*. He is a member of the IET and serves on the Executive Committee of the British Machine Vision Association.

Xianghua Xie received the M.Sc. (with commendation) and Ph.D. degrees in computer science from the University of Bristol, Bristol, U.K., in 2002 and 2006, respectively. He is currently an Associate Professor with the Department of Computer Science, Swansea University, Swansea, U.K. He was a Research Associate with the Department of Computer Science, University of Bristol, between 2006 and 2007. His current research interests are medical imaging, video analysis, texture analysis, image segmentation, surface inspection, deformable models, and human pose estimation and tracking. He is a member of BMVA.

Mark C. K. Hamilton received the Medical degree from the University of Dundee in 1991, Membership of the Royal College of Physicians in 1995 and Fellowship of the Royal College of Radiologists in 2000. He was a Cardiac Radiology Fellow in Auckland, New Zealand from 2003 to 2005. Since then, he has been a Consultant Cardiac Radiologist at United Hospitals Bristol, and clinical lead in cardiac magnetic resonance imaging. His research collaborates with the Bristol Cardiovascular NIHR Biomedical Research Unit and the Department of Computer Science, Bristol University. His research interests in computed tomography (in particular protocol optimisation) and cardiac magnetic resonance imaging (in particular exercise physiology in health and disease).