

● *Original Contribution*

## UNSUPERVISED IMAGE CLASSIFICATION OF MEDICAL ULTRASOUND DATA BY MULTIREOLUTION ELASTIC REGISTRATION

SCHLOMO V. ASCHKENASY,<sup>\*†</sup> CHRISTIAN JANSEN,<sup>\*†</sup> REMO OSTERWALDER,<sup>†</sup> ANDRÉ LINKA,<sup>†</sup>  
MICHAEL UNSER,<sup>‡</sup> STEPHAN MARSCH,<sup>\*</sup> and PATRICK HUNZIKER<sup>\*</sup>

<sup>\*</sup>Medical Intensive Care Unit, University Hospital, Basel; <sup>†</sup>Department of Cardiology, University Hospital, Basel; and <sup>‡</sup>Biomedical Imaging Group, Swiss Federal Institute of Technology, Lausanne, Switzerland

(Received 4 April 2005, revised 8 March 2006, in final form 16 March 2006)

**Abstract**—Thousands of medical images are saved in databases every day and the need for algorithms able to handle such data in an unsupervised manner is steadily increasing. The classification of ultrasound images is an outstandingly difficult task, due to the high noise level of these images. We present a detailed description of an algorithm based on multiscale elastic registration capable of unsupervised, landmark-free classification of cardiac ultrasound images into their respective views (apical four chamber, two chamber, parasternal long axis and short axis views). We validated the algorithm with 90 unselected, consecutive echocardiographic images recorded during daily clinical work. When the two visually very similar apical views (four chamber and two chamber) are combined into one class, we obtained a 93.0% correct classification ( $\chi^2 = 123.8$ ,  $p < 0.0001$ , cross-validation 93.0%;  $\chi^2 = 131.1$ ,  $p < 0.0001$ ). Classification into the 4 classes reached a 90.0% correct classification ( $\chi^2 = 205.4$ ,  $p < 0.0001$ , cross-validation 82.2%;  $\chi^2 = 165.9$ ,  $p < 0.0001$ ). (E-mail: hunzikerp@uhbs.ch) © 2006 World Federation for Ultrasound in Medicine & Biology.

**Key Words:** Multiresolution, Multiscale, Registration, Warp, Image matching, Database retrieval, View classification.

### INTRODUCTION

#### *Motivation*

The automatic identification of image content in medical images is an important preprocessing step in computer-assisted diagnosis, content-based image retrieval and picture archival systems. Most database systems used in clinical environments today request a textural query, which implies that images are manually classified. This task is tedious, error prone and often not suitable for legacy images. For high-quality chest radiographs, an algorithm capable of classifying radiographs into frontal and lateral views has been described (Lehmann et al. 2003) and there is ongoing work in computed tomography and magnetic resonance imaging (Glatard et al. 2004). The classification of ultrasound (US) images is particularly difficult, because US images have high speckle-noise content, show nonconstant intensities even for comparable structures and typically have discontin-

uous structural boundaries, setting US data apart from other image modalities used in medicine.

Echocardiography is a typical application of US imaging and is the most frequently used tool in clinical practice for the evaluation of human heart morphology and function. During the examination, still images and image loops are recorded from several transducer locations on the chest, leading to characteristic views that show the same organ from different but distinct viewpoints. Figure 1 shows representative examples of four standard views (apical four chamber- and two chamber views, parasternal long axis and short axis views). In this paper, we present an algorithm capable of an unsupervised landmark-free classification of such images into their respective views.

#### *Approach*

Simple template matching is difficult to apply for echocardiographic image classification. Differences in heart size, shape and image intensities, combined with a multitude of disease processes (each characterized by alteration of different features of heart anatomy) would require a huge template library and would render com-

Address correspondence to: PD Dr. P. Hunziker, Universitätsspital Basel, Petersgraben 4, 4031 Basel, Switzerland. E-mail: hunzikerp@uhbs.ch




Fig. 1. Typical representations of the four view classes (preprocessed with local and global normalization). The four and two chamber views (A4 and A2) are recorded from the hearts apex and therefore called apical views. The long axis and short axis views (LAX and SAX) are recorded from a position left of the sternum and therefore called parasternal views.

putational cost prohibitive. We therefore chose an adaptive approach that uses multiscale elastic registration to match the unknown sample image onto known templates. This computer vision approach allows the identification of images by the use of only one or very few templates per class.

By matching a sample image against the template library, we obtain for each sample image a deformation map and warped sample images corresponding to each of the templates. Both deformation energy and the similarity of the warped and the reference images can then be used to classify the sample image. The smaller the deformation energy needed for optimal matching (calculated from the deformation map) and the more similar the resulting image (difference between template and warped sample), the higher is the probability that the sample image and the template image belong to the same class. Figure 2 gives an overview of the classification process.

To achieve this, a multiscale approach for both the images and the displacement map was chosen; this cir-

cumvents the problem of local minima during registration and reduces computational cost for the optimization process. A gradient descent optimizer was used for registration. To ensure a sufficient degree of spatial coherence, which is implied by the biophysics of the heart, the allowed deformation was limited to that fitting into a continuous B-spline-based model.

The use of splines throughout all aspects of the algorithm, namely the model, scaling, warping and interpolation, has several advantages, including their good approximation properties, the easy and fast transformation from the discrete pixel-domain to the continuous spline domain that allows spline interpolation, warping and multidimensional scaling using standard filtering techniques.

Models using cubic splines are physically plausible (Ahlberg et al. 1967) and can represent all affine transformations. Moreover, splines are scaleable in the sense that any coarse level deformation can be represented at a finer scale without loss of information, given an integer ratio between scales. An excellent overview on the sub-

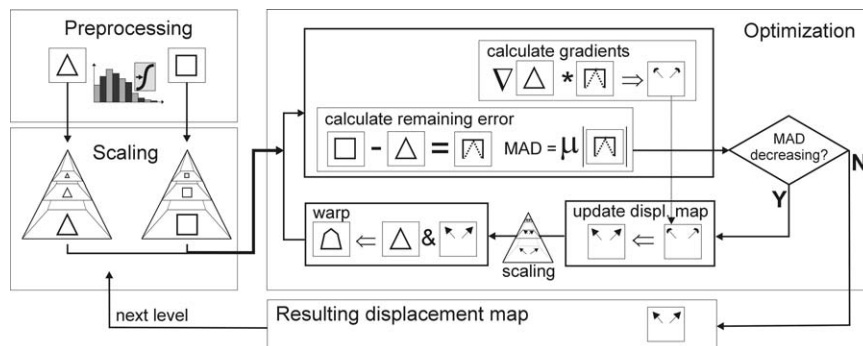


Fig. 2. Graphical representation of the registration algorithm Images are preprocessed before they are placed in the pyramid. Scaled images are fed into the optimization loop and adapt the displacement map until the MAD fails to decrease. Once the optimum on this level is reached, the algorithm proceeds to the next pyramid level with higher resolving images or displacement maps in predefined order.

ject of splines in medical imaging can be found online at <http://bigwww.epfl.ch/publications/unser0002.html>.

## METHODS

### Registration Algorithm

We used a multidimensional, multiscale, intensity-based image registration algorithm based on parametric spline models to describe both the image and the deformation map. Working in the spline domain offers continuous data and deformation models and is attractive because of the availability of fast algorithms and analytically defined partial derivatives and integrals.

The input images are given as two two-dimensional (2-D) discrete signals, the template image  $f_{Tj}[i]$  and the sample image  $f_S[i]$  to be classified. The index is  $i \in IC\mathbb{Z}^2$ , where  $I$  is a 2-D discrete interval representing the set of all pixel coordinates in the image. The images were transferred into continuous spline representations using a third-order direct B-spline transform filter as described in (Unser *et al.* 1993a, 1993b).

$$f_{Tj}[i] \Rightarrow f_{Tj}(i) \quad f_S[i] \Rightarrow f_S(i) \quad (1)$$

The local deformation  $u$  was described as a continuous B-spline model of displacement vectors

$$\mathbf{u}(i) = \begin{pmatrix} u_x(i) \\ u_y(i) \end{pmatrix} \quad (2)$$

each built from a lower number of coefficients (coarser resolution) at each pyramid level compared with the image data. This implicit spatial constraint enforces smoothness of the solution in space, solves the problem of under-determination and helps to reduce the impact of noise.

The deformation  $u$  transforms the sample image  $f_S(i)$  into an image  $g$  that is to become as similar as possible to the template image  $f_{Tj}(i)$  during the optimization process.

$$i = (i_x, i_y) \quad g(i) = f_S(i_x - u_x, i_y - u_y) \sim f_{Tj}(i) \quad (3)$$

$g(i) - f_{Tj}(i)$  corresponds to the difference in intensity and eqn (3) can therefore be considered to be an optimization problem, where the mean absolute deviation (MAD) (eqn (4)) of the two images is minimized

$$MAD = \frac{1}{\|I\|} \sum_{i \in I} |g(i) - f_{Tj}(i)|. \quad (4)$$

We used the MAD and not the mean-squared error to render our system less sensitive to the shot-noise typically found in US images.

During the optimization process driven by a gradient descent algorithm (Press *et al.* 2002) along the local

intensity gradients,  $f_S$  was progressively warped using the deformation models  $u_x$  and  $u_y$ . The deformation model was in turn incrementally updated according to the local (smoothed) intensity gradients in  $g$ . Optimization on each level was terminated as soon as MAD increased or when two successive optimization loops did not decrease the MAD by a prespecified level (we used  $10^{-3}$  gray-levels per pixel).

The algorithm is a variation of the elastic registration procedure proposed by several groups of researchers (Kybic and Unser 2003; Mattes *et al.* 2003; Musse *et al.* 2004).

### Multiresolution Pyramids

We worked with a multiscale pyramid, to allow large displacements without violating the Nyquist theorem, to prevent the algorithm from being trapped in local minima and to speed up computation by working on smaller datasets. Down- and up-sampling of images and deformation maps in pyramids was done using fast least-square spline approximation using the digital filtering approach described in (Unser *et al.* 1993a, 1993b). The down-sampling filter decimates each axis by a factor of two, leading to a data reduction of 1:2<sup>2</sup> per pyramid level.

### Image Warping

Cubic B-spline interpolation of the sample images according to the calculated deformation map was done progressively to warp the sample image until “best fit” was achieved. This type of interpolation minimizes loss as a result of resampling (Unser *et al.* 1993a, 1993b).

### Classification

Each sample image thus yields one elastically registered/warped image for each template. Classification was done using linear discriminant analysis, whereby the classifier was constructed by an optimal combination of (a) on a measure of similarity, and (b) a measure of the effort needed for elastic registration. As measure of similarity, the resulting MAD (eqn (4)) after image registration was used. As measure of effort needed for elastic registration, we used a linear combination of parameters that were calculated from the resulting displacement map: they were displacement, strain and variation of strain.

The first parameter is the displacement for optimal registration. To compensate for zero order movement as a result of simple shifting, the mean overall displacement is subtracted.

$$u_{\max} = \max_{i \in I} \|\mathbf{u}(i) - \mu_u\|$$

$$\mu_u = \left[ \frac{1}{\|I\|} \sum_{i \in I} u_x(i), \frac{1}{\|I\|} \sum_{i \in I} u_y(i) \right]^T. \quad (5)$$

Table 1. Pseudo code of the algorithm.

---

```

BEGIN
  NormalizeGrayscale;
  CreatePyramid(Sample0);
  CreatePyramid(Model);
  FOR all templates DO
    Sample: = Sample0;
    CreatePyramid(Template);
    FOR level: = coarse TO fine DO
      WHILE Difference(Sample, Template) > threshold DO
        Gradient: = FindGradient(Sample, Template);
        Model: = UpdateModel(Model, Gradient);
        Sample: = WarpImage(Sample0, Model);
      END;
    END;
    Parameters[template]: = CalcParameters(Difference, Model)
  END;
  View: = LinearDiscriminantAnalysis(Parameters);
END;

```

---

The other two parameters are based on the variation of the displacement vectors denoting the amount of strain needed to make the images match. They are calculated from the partial derivatives of the deformation model  $u$ . The amount of strain (eqn (6))  $s$  in  $u$

$$s_{\max} = \max_{i \in I} \|s(i)\| \quad s = \left[ \frac{\partial u}{\partial x}, \frac{\partial u}{\partial y} \right]^T \quad (6)$$

and the standard deviation  $\rho_s$  of  $s(i)$  (eqn (7))

$$\sigma_s = \frac{1}{\|I\|} \sum_{i=1} \|s(i) - \mu_s\|$$

$$\mu_s = \left[ \frac{1}{\|I\|} \sum_{i \in I} s_x(i), \frac{1}{\|I\|} \sum_{i \in I} s_y(i) \right]. \quad (7)$$

The rationale behind these parameters is that we expect the need of a small displacement (eqn (5)), minimal deformation or strain (eqn (6)) of maximal homogeneity and smoothness (eqn (7)) to match two images of the same class.

Actual classification was done by linear discriminant analysis. The three canonical functions and their respective eigen values were determined by canonical correlation analysis. To eliminate possible overfitting of the strategy, which could result in apparently perfect fitting of the training data at the expense of worse fitting of new samples, cross-validation using a “leave-one-out” strategy was done; the sample to classify was not used for constructing the classifier function.

Table 1 shows an overview of the algorithm using pseudo-code, whereas Fig. 2 gives graphical overview of the algorithm.

### Templates

We found that templates with high echo quality, in particular, with highly visible ventricular walls, per-

formed best during classification because this increases the probability for two correlating image areas to overlap. Templates were defined by expert designation of “prototypical” images.

### Image Preprocessing

To achieve optimal classification of the clinical image set, where acquisition parameters are typically chosen by the subjective preferences of the echocardiographer, preprocessing of the images to achieve a similar magnification and intensity distribution was important. We utilized the information on physical dimensions and size stored in digital US images to adjust the templates to match the size calibration of each given sample image before starting registration. Then, image intensities were normalized both locally and globally. Local normalization was done by equalization of local mean and variance within an image (eqn (8)) using a sliding window technique (window size 20% and 40% of the image width)

$$n(i) = \frac{f(i) - \mu_f(i)}{\sigma_f(i)}. \quad (8)$$

Global normalization was achieved by equalizing the cumulated histogram of the locally normalized image to the error function (*erf*) as described in Bosch et al. (2002).

Masks were created for the template and the test image and were combined by logical AND to build a single mask that removes all but the area containing image data in both images. Electrocardiographic signal curves sometimes overlaid on nonmasked image parts were removed by interpolation.

### Patients

The 90 echocardiographs used to test and validate the algorithm were consecutive nonselected image acquisitions recorded by different sonographers during daily clinical work in a cardiology outpatient clinic using a Hewlett-Packard Sonos 5500 US device. None of the patients had congenital heart defects. The mean age was  $48 \pm 12$  years. The sole quality-criterion for the images to be included was the fact that an experienced cardiologist was able to classify the image into one of the four classes used

## RESULTS

### Elastic Registration

Our algorithm uses a progressively finer displacement map to adapt the template to the sample image (multiscale approach) compared with other algorithms that start with a fine displacement map in the first place (single scale approach). In spite of similar result-images, multiscale registration prevents trapping in local minima




Fig. 3. Elastic registration using a single vs. a multiresolution approach. The image to be classified (a) and the template image (b) are registered. When multiscale registration is used, a good fit (a2) and a coherent motion map (b2) are achieved. Single-scale registration on the contrary leads to a chaotic, incoherent motion map (b1) and a bad fit (a1). The vectors in (b1) were multiplied by 10 to improve visualization.

and therefore results in a coherent displacement map, whereas single scale registration results in a chaotic deformation map (Fig. 3).

Optimization by increasing the resolution in the deformation map during the early iteration steps before the image resolution was increased was found to improve both the accuracy and the resulting image quality.

Figure 4 shows that the most significant part of the displacement vectors can be calculated on very low resolutions. A later increase in image resolution will lead to a relatively small but significant improvement in the quality of the final image.

#### Classification

Classification results of the 90 echocardiographic images to four standard echocardiographic view tem-

plates (parasternal short-axis, parasternal long-axis, apical four-chamber and apical two-chamber views) are given in Table 2. Using MAD, the largest displacement of  $u$ , the largest strain and its standard deviation, we were able to classify 90.0% of the images to one of four classes ( $\chi^2 = 205.4$ ,  $p < 0.0001$ ).

To verify this result, we used a “leave-one-out” strategy and classified each of the 90 images using a system trained by the remaining 89 samples only. This cross-validation yielded an 82.2% correct classification ( $\chi^2 = 165.9$ ,  $p < 0.0001$ ), indicating robust classification capabilities to new samples not in the training set.

When the two of the classes that are visually very similar (four-chamber and two-chamber views) were combined into one class, this yielded a 93% correct classification for both the original ( $\chi^2 = 123.8$ ,  $p < 0.0001$ ) and the cross-validated ( $\chi^2 = 131.1$ ,  $p < 0.0001$ ) classification strategies. These results are presented in Table 3.

Using MAD only as a classification criterion led to only 74.4% correctly classified images, and “leave-one-out” cross-validation then yielded 72.2% ( $\chi^2 = 48.8$ ,  $p < 0.0001$ ) correct classification.

#### Speed

The time needed to register an image elastically varies as a function of the initial similarity of the images. On an Intel Pentium 4 2.8-GHz standard personal computer with 1 GB RAM running on Microsoft Windows XP, median classification duration was 3.4 s (interquartile width 2.4 s to 4.5 s) for an image size of  $720 \times 512$  pixels.

## DISCUSSION

US images with their high speckle-noise content, nonconstant intensities and discontinuous boundaries pose particular difficulties for image processing algorithms. Multiscale elastic registration using continuous spline models of the image and the deformation is here shown to achieve fast and accurate unsupervised classification of unselected echocardiographic images acquired during routine clinical examinations.

Several aspects of the presented algorithm deserve to be mentioned. We chose a multiresolution approach for both the images and the displacement maps. This approach has several advantages (Sühling *et al.* 2004): multiresolution in the deformation map improves robustness because it reduces sensitivity of the algorithm to be trapped in local minima. In addition to this, the smoothness implicit in a multiresolution model enforces coherent deformation in adjacent image areas that are required by the biophysics of tissue; thus, the severe noise problem inherent in US images might be overcome. In a typical classification, image size


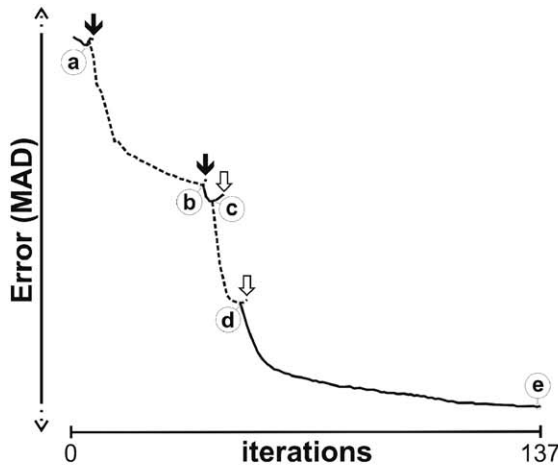


Fig. 4. Improvement of fit during optimization The graph shows a typical course of the MAD during optimization. The arrows denote a step on the deformation map pyramid ( $\downarrow$ ) resulting in a higher deformation resolution and a step in the image pyramid ( $\downarrow$ ) resulting in a higher image resolution. The images (a) to (e) denote the intermediate images after each of the optimization steps, as marked in their respective positions on the graph.

was reduced by ( $\sim 28:1$ ) and the displacement map by ( $\sim 200:1$ ) for the initial step.

Because optical flow is mathematically underdetermined by a factor of two, we used a displacement map of one-quarter of the image size, thereby reducing the system by four degrees of freedom and constraining the algorithm to a smooth solution.

Table 2. Classification results using four template classes. Bold numbers of the diagonal denote the correctly classified samples, 90% for the upper part entitled “Original” and 82.2% for the bottom part entitled “Cross validated”

	Class	Predicted image type				Total
		A4	A2	LAX	SAX	
<b>Original</b>	<b>A4</b>	<b>21</b>	0	2	1	24
	<b>A2</b>	3	<b>21</b>	0	1	25
	<b>LAX</b>	0	0	<b>24</b>	0	24
	<b>SAX</b>	2	0	0	<b>15</b>	17
<b>Cross-validated</b>	<b>A4</b>	<b>20</b>	1	2	1	24
	<b>A2</b>	6	<b>17</b>	0	2	25
	<b>LAX</b>	0	0	<b>24</b>	0	24
	<b>SAX</b>	4	0	0	<b>13</b>	17

(Abbreviations: A4: apical four chamber view; A2: apical two chamber view; LAX: long axis view; SAX: short axis view).

Figure 3 shows the result of a single scale (image a1, map b1) and multiscale (image a2, map b2) registration. Although a single scale approach leads to rather incoherent rearrangement of pixels, a multiscale approach overcomes the problems caused by the nature of US images and leads to an excellent match.

Going down the resolution pyramid to the finest level yielded near-perfect matching at the expense of very complex deformation maps. However, for image classification purposes, fewer scales can be used when

Table 3. Classification results using three template classes

	Class	Predicted image type			Total
		A	LAX	SAX	
<b>Original</b>	<b>A</b>	<b>45</b>	2	2	49
	<b>LAX</b>	0	<b>24</b>	0	24
	<b>SAX</b>	2	0	<b>15</b>	17
<b>Cross-validated</b>	<b>A</b>	<b>44</b>	2	3	49
	<b>LAX</b>	0	<b>24</b>	0	24
	<b>SAX</b>	4	0	<b>13</b>	17

Bold numbers of the diagonal denote the correctly classified samples (93% in both cases). (Abbreviations: A: apical view (includes four- and two-chamber views); LAX: long axis view; SAX: short axis view).

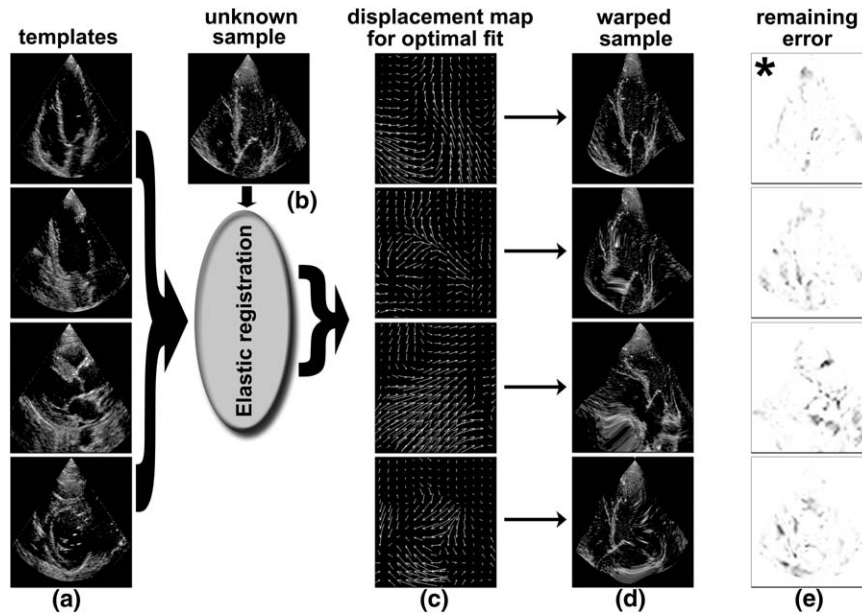


Fig. 5. Classification of one sample into one of four templates (a) Four templates; (b) the sample image; (c) four displacement maps resulting from the algorithm; (d) template images warped using the displacement maps; and (e) difference-images of the warped templates and the sample image. \* denotes correctly identified class. The algorithm registers four templates with a sample image resulting in four displacement maps. The maps can be used to warp the sample image into an image that is as similar as possible to each of the templates. Note the large deformation necessary to match a sample to a nonfitting template. The difference (MAD) between the four warped sample images and the template images yields four error images. The MAD of the difference images as well as the parameters derived from the displacement map (see text) are then used in linear discriminant analysis to find the class membership.

overall shape is deemed to be more important than perfect fitting of minute texture details in nonidentical hearts. The high-resolution fitting approach renders the classification algorithm more dependent on the deformation information, whereas in a coarser-scale approach, the classification algorithm relies more on the final fitting information. We found that stopping the algorithm at a (linear) resolution of 1:8 for the image data yields classification results that were not inferior to an approach going down to the ultimate resolution, but reduced the number of iterations from 100 to 200 in the fine resolution case to 10 to 25.

As to speed, the current version, although not using optimized libraries, allows reliable classification of more than 1000 images per hour on a standard desktop computer and appears thus to be suitable for practical use in a hospital environment, even in the absence of high-speed computing facilities. Based on experiments with some optimized libraries, we expect the potential of a five- to 10-fold speed-up of the algorithm on the same hardware.

In spite of the complex nature of US images and the low image quality of the unselected images, we were able to classify 90% of the images. “Leave-one-out” cross-validation yielded an 82.2% correct classification. This result improved significantly (93% for both) when

the two visually very similar classes (apical four-chamber and apical two-chamber views) were combined into one class (“apical views”). Figure 1 demonstrates clearly that large image parts of these two classes look very similar. In some cases, their discrimination is difficult even for specialists. Ongoing work shows promising results when images classified as being “apical view” were separated in a second step by using only the lower part of the image in a second classification step.

#### Limitations

A limitation of our algorithm is the need for templates that are specific for each of the echocardiographic views to classify. Although these views are clearly standardized, they still vary significantly, depending on the anatomy of the patient and the skill of the operator. This problem of “missing gold standard” could be overcome by using several templates for each of the views to classify or a system that dynamically adds new template classes if a sample does not fit. Our ongoing work studies the use of classification results for building better class-templates in a “bootstrap” fashion.

We used gradient descent for an optimization method, although more complex algorithms are available for multidimensional optimization. Because the classification results shown could be achieved by using approx-

imately 10 to 25 iterations per sample and template, the speed trade-off per iteration caused by more complex multidimensional optimization algorithms for registration could cost as much computation time as is gained by the reduction of necessary iterations. Interestingly, this trade-off has been shown for the Marquard-Levenberg algorithm in an approach studying affine registration of multidimensional datasets in [Kybic and Unser \(2003\)](#), where the authors confirmed that the strength of a refined optimizer lies predominantly in the final high-resolution steps in such multiresolution approaches.

#### *Related Work*

[Lehmann et al. \(2003\)](#) propose an algorithm to classify chest radiographs into their respective views. They reduce these images to  $32 \times 32$  pixels, independent of the original size, and tangent distance was then used for their classification. This algorithm is not well suited for echocardiographic images, mainly because some echocardiographic views look similar in terms of brightness distribution because echo images often do not have large areas of similar brightness (as the lung-in-chest images) and because in echocardiograms the bright but variable lung tissue at the image borders would tend to dominate such an algorithm.

In some face-recognition algorithms ([Blanz and Vetter 2003](#); [Zhao et al. 2003](#)), a well-studied field where image classification is the goal as well, optimal positioning of a face before the actual classification task takes place is an important preprocessing step. Elastic registration similar to the one we used can be used for this purpose, although there are differences in the practical problems posed in face recognition compared with US. However, the success of algorithms based on multiresolution ([Raducanu et al. 2001](#)), elastic matching ([Blanz and Vetter 2003](#)) and template libraries ([Blanz and Vetter 2003](#); [Hallinan 1991](#)) in face recognition certainly underscore the usefulness of these ingredients in classification approaches in fields such as medical imaging.

## CONCLUSION

We present a multiscale elastic registration algorithm based on a continuous model of both images and deformation maps and used it successfully in unsupervised classification of non selected cardiac ultrasound images acquired during daily clinical practice by means of a template library. [Figure 5](#)

## REFERENCES

- Ahlberg JH, Nilson NE, Walsh JL. The theory of splines and their applications. New York: Academic Press, 1967. pp. xi, 284.
- Blanz V, Vetter T. Face recognition based on fitting a 3D morphable model. *IEEE Trans Pattern Anal Machine Intel* 2003;25:1063–1074.
- Bosch JG, Mitchell SC, Lelieveldt BPF, Nijland F, Sonka M, Reiber JHC. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE Trans Med Imag* 2002;21:1374–1383.
- Glatard T, Montagnat J, Magnin I. Texture based medical image indexing and retrieval: Application to cardiac imaging. *Proceedings of the 6th ACM SIGMM international workshop on Multimedia information retrieval 2004*;135–142.
- Hallinan PW. Recognizing human eyes. *G 1991*; SPIE Proceedings 1570;214–226.
- Kybic J, Unser M. Fast parametric elastic image registration. *IEEE Transactions on Image Processing* 2003;12:1427–1442.
- Lehmann TM, Güld O, Keyzers D, Henning S, Kohnen M, Wein B. Determining the view of chest radiographs. *J Digital Imag* 2003; 16:280–291.
- Mattes D, Haynor D, Vesselle H, Lewellen T, Eubank W. PET-CT image registration in the chest using free-form deformations. *IEEE Trans Med Imag* 2003;22:120–128.
- Musse O, Heitz F, Armspach JP. Fast deformable matching of 3D images over multiscale nested subspaces. Application to atlas-based MRI segmentation. *Pattern Recogn* 2004;36:1881–1899.
- Press, WH, Teukolsky, SA, Vetterling, WT, Flannery, BP. *Minimization or maximization of functions in numerical recipes in C*. Cambridge, MA: University Press, 2002.
- Raducanu B, Grana M, Albizuri FX, d'Anjou A. Face localization based on the morphological multiscale fingerprints. *Pattern Recogn Lett* 2001;22:359–371.
- Sühling M, Jansen C, Arigovindan M, et al. Multiscale motion mapping—A novel computer vision technique for quantitative, objective echocardiographic motion measurement independent of Doppler: First clinical description and validation. *Circulation* 2004;110: 3093–3099.
- Zhao W, Chellappa R, Phillips J, Rosenfeld A. Face recognition: A literature survey. *ACM Comput Surv* 2003;35:399–458.