

REDUCING THE TRAINING SET USING SEMI-SUPERVISED SELF-TRAINING ALGORITHM FOR SEGMENTING THE LEFT VENTRICLE IN ULTRASOUND IMAGES

Jacinto C. Nascimento

Gustavo Carneiro

Instituto de Sistemas e Robótica, Instituto Superior Técnico, 1049-001 Lisboa, Portugal

ABSTRACT

Statistical pattern recognition models are one of the core research topics in the segmentation of the left ventricle of the heart from ultrasound data. The underlying statistical model usually relies on a strong prior for the shape and appearance of the left ventricle whose parameters can be learned using a manually segmented data set. Unfortunately, this is usually quite complex, requiring a large number of parameters that can be robustly learned only if the training set is sufficiently large. The difficulty in obtaining large training sets is currently a major roadblock for the further exploration of statistical models in medical image analysis. In this paper, we present a novel semi-supervised self-training model that reduces the need of large training sets for estimating the parameters of statistical models. This model is initially trained with a small set of manually segmented images, and for each new test sequence, the system re-estimates the model parameters incrementally without any further manual intervention. We show that state-of-the-art segmentation results can be achieved with training sets containing 50 annotated examples.

1. INTRODUCTION

The segmentation of the left ventricle (LV) of the heart is a major topic of research in the area of medical image analysis. In practice, the automatic LV segmentation represents an important tool in clinical settings due to the following reasons: 1) it can increase patient throughput; and 2) it can reduce inter-user variation in the LV delineation procedure.

One of the main techniques employed to solve this problem is based on statistical pattern recognition approaches [1–3], which model the LV appearance and shape using a set of manually annotated images (i.e. the training set). This model is then used for building a classifier that detects and segments the LV from ultrasound data. The models used in statistical pattern recognition approaches usually contains from hundreds to thousands of parameters, which can only be robustly estimated with large training sets, a fact usually known as the *curse of dimensionality*. Actually, it is not uncommon that systems based on such approaches need in the order of thousands of images, coming from several different sources and annotated by different clinicians [2,4]. However, the acquisition of such large training sets is a formidable task, demanding an extremely large amount of time from clinicians. The fact that the majority of researchers working in this field cannot have access to such large training sets results in insufficient investigation of statistical pattern recognition models.

The machine learning and computer vision communities have faced similar issues over the last few years, which resulted in the development of semi-supervised learning techniques [5]. These

This work was supported by the FCT (ISR/IST plurianual funding) through the PIDDAC Program funds and Project HEARTRACK (PTDC/EEA-CRO/103462/2008). This work was partially funded by EU Project IMASEG3D (PIIF-GA-2009-236173).




Fig. 1. Semi-supervised learning. The graph on the left shows the classification problem where only a small subset of the samples are labeled. The graph on the right displays the result of semi-supervised learning, where $P(o|x)$ is the probability of class o given point x .

learning methods use both annotated and unannotated training data to estimate the parameters of statistical models. The main assumption is that samples belonging to the same class tend to cluster together in the input space, and if a few annotated examples are given, we can associate unannotated samples of the cluster with the label of annotated samples in that same cluster, as shown in Fig. 1.

In this paper, we introduce a semi-supervised learning approach that initially trains a statistical model of LV shape and appearance using a small training set. This training set is used to build an initial classifier that detects and segments the LV from ultrasound data. Given a previously unseen test sequence, the system uses the classifier to detect positive and negative hypotheses in each frame. These hypotheses are added to the training set if the detection confidence is above a certain threshold, and the model is then re-trained with the updated training set. This algorithm is an instance of the semi-supervised self-training learning approach. Our approach is innovative in the sense that we do not rely on a fixed set of unannotated images, which is a common assumption that helps the self-training process.

2. SELF-TRAINING

Assume that $\mathbf{x} \in \mathbb{R}^D$ denotes the feature vector representing the data (e.g., image), $\mathbf{y} \in \mathbb{R}^N$ represents the annotation (e.g., manual LV segmentation), the hypothesis confidence is measured with the posterior classifier $p(\mathbf{y}|\mathbf{x})$, and the data density is represented by $p(\mathbf{x})$. Self-training methods assumes that $p(\mathbf{y}|\mathbf{x})$ and $p(\mathbf{x})$ share parameters in order to train the classifier $p(\mathbf{y}|\mathbf{x})$ such that annotation transitions can happen only at locations where the density of $p(\mathbf{x})$ is low [5].

For the derivation of our algorithm, consider that the set of training images is represented by \mathcal{X} , and \mathcal{Y} denotes the respective set of manual annotations. The goal of the self-training is to estimate the parameters θ of the classifier $p(\mathbf{y}|\mathbf{x}, \theta)$ using the annotated training set $\{\mathcal{X}, \mathcal{Y}\}$ and a set of unannotated images $\{\tilde{\mathbf{x}}_i\}_{i=1..K}$ along with the probabilities of producing the respective annotations $\{\tilde{\mathbf{y}}_i\}_{i=1..K}$ given by $p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta)$. This is summa-

rized as [6]:

$$\begin{aligned} \theta^* &= \arg \max_{\theta} P(\mathcal{Y}|\mathcal{X}, \theta) \\ &\propto \arg \max_{\theta} \log \sum_i f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \frac{p(\mathcal{Y}, \tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \mathcal{X}, \theta)}{f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i)}, \end{aligned} \quad (1)$$

where $f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) : \mathbb{R}^D \times \mathbb{R}^N \rightarrow \mathbb{R}$ has the constraints $\sum_i f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) = 1$ and $f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \geq 0$. Using Jensen's inequality, we can find the following lower bound to the objective function (1):

$$\begin{aligned} \underbrace{\sum_i f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \log \frac{p(\mathcal{Y}, \tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \mathcal{X}, \theta)}{f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i)}}_{\text{Lower Bound}} &\leq \\ \log \sum_i f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \frac{p(\mathcal{Y}, \tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \mathcal{X}, \theta)}{f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i)} & \end{aligned} \quad (2)$$

This lower bound is easier to maximize than the original objective function (1). Therefore, we solve the following optimization problem:

$$\theta^* = \arg \max_{\theta} \sum_i f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \log \frac{p(\mathcal{Y}, \tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \mathcal{X}, \theta)}{f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i)} \quad (3)$$

subject to, $\sum_i f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) = 1$, $f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \geq 0$. Taking the derivative of the Lagrangian with respect to $f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i)$, we find:

$$f(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) = p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta). \quad (4)$$

Hence, we can formulate an iterative algorithm comprising the following expectation (E) and maximization (M) steps:

- **E-step:**

$$f^{(t)}(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) = p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta^{(t-1)}) \quad (5)$$

- **M-step:**

$$\theta^{(t)} = \arg \max_{\theta} \mathbb{E}_{f^{(t)}(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i)} \left[\log p(\mathcal{Y}, \tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \mathcal{X}, \theta) \right], \quad (6)$$

where the superscript (t) indicates the iteration index.

Therefore, we propose an iterative on-line EM algorithm (see Alg. 1), where the goal is to maximize (with respect to θ) and generalize (in the data space \mathbf{x}) the model $p(\mathbf{y}|\mathbf{x}, \theta)$ with the constraint that there are no transitions of $p(\mathbf{y}|\mathbf{x}, \theta)$ on high density regions of $p(\mathbf{x})$. Both the generalization goal and the constraint are achieved by incrementally incorporating in the training set examples $(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)$ that produced $p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta) \geq \gamma$, where $\gamma > 0$ is a free variable. We stress that, similar self-training algorithms use a fixed size set of unannotated images, which is different from our problem since the test images are dynamically provided when a new sequence is given to the system to process.

There are three issues to discuss in Alg 1. The first point is that $\theta^{(0)}$ is obtained from the maximization of $p(\mathcal{Y}|\mathcal{X}, \theta)$ using the annotated training data only. Hereafter, we denote the training process to estimate $\theta^{(0)}$ as *supervised*, and the training for $\theta^{(t)}$ for $t > 0$ as *semi-supervised* (in the sense that unannotated samples will be incorporated in the learning scheme). The second is that we select the samples to be included in the training set by sampling a Gaussian mixture model (7) and taking $\tilde{\mathbf{x}}_i$ (with annotation $\tilde{\mathbf{y}}_i$) with probability $p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta^{(t-1)}) \times \mathcal{N}(\tilde{\mathbf{y}}_i, \Sigma)$, where $\mathcal{N}(\mu, \Sigma)$ is the Gaussian probability density function with mean μ and covariance Σ (we set Σ to be $10^{-3} \times \mathbf{I}$, with \mathbf{I} the identity matrix). In this work, the number of samples drawn from (7) is the same as the size of the training set $|\{\mathcal{X}, \mathcal{Y}\}|$. The third point, which we provide more details in the next section, is that the classifier $p(\mathbf{y}|\mathbf{x}, \theta)$ is based on deep learning methods. Note that in Alg. 1, T denotes the maximum number of iterations, which is set to 100.

Algorithm 1 Iterative on-line EM

for $t = 1:T$ do
E-STEP: Sample

$$(\tilde{\mathcal{Y}}, \tilde{\mathcal{X}}) \sim \sum_i p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta^{(t-1)}) \times \mathcal{N}(\tilde{\mathbf{y}}_i, \Sigma) \quad (7)$$

$\tilde{\mathcal{Y}} = \mathcal{Y} \cup \tilde{\mathcal{Y}}, \tilde{\mathcal{X}} = \mathcal{X} \cup \tilde{\mathcal{X}}$

M-STEP:
maximize $\theta^{(t)} p(\tilde{\mathcal{Y}}|\tilde{\mathcal{X}}, \theta)$

subject to: (for all $(\tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_i) \in (\tilde{\mathcal{Y}}, \tilde{\mathcal{X}})$)

$$\begin{aligned} p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta^{(t-1)}) &\geq \gamma \\ \sum_i p(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i, \theta^{(t-1)}) &= 1 \end{aligned} \quad (8)$$

if $\|\theta^{(t)} - \theta^{(t-1)}\|_2 \leq \epsilon$ then STOP iterations.
end for




Fig. 2. Original training image (top left) with the manual LV segmentation in yellow line and star markers (top middle) with the rectangular patch representing a canonical coordinate system for the segmentation markers. The top-right image shows the reference patch with the base and apical points highlighted and located at their canonical locations within the patch (these points are used to define the rigid transform of the patch). The images on the second and third rows display several positive and negative patches (respectively) used to train the rigid classifier.

3. SEGMENTING THE LEFT VENTRICLE USING DEEP LEARNING METHODS

The classifier $p(\mathbf{y}|\mathbf{x}, \theta)$ is based on deep neural networks [7], which is a type of deep learning classifier. Deep neural networks have been recently explored by Carneiro et al. [1], who showed that this classifier can achieve state-of-the-art LV segmentation results with 400 annotated training images.

Consider that $\mathbf{y} = [s_j]_{j=1..N}$ represents the vector of keypoints $s_j \in \mathbb{R}^2$ for the LV segmentation of an ultrasound image I . The annotated training set is denoted by $\mathcal{D} = \{(I, \vartheta, \mathbf{y})\}_{i=1..M}$, with LV images I_i , the respective manual annotation \mathbf{y}_i and the parameters of a rigid transformation $\vartheta_i \in \mathbb{R}^5$ (position $\mathbf{p} \in \mathbb{R}^2$, orientation $\xi \in [-\pi, \pi]$, and scale $\sigma \in \mathbb{R}^2$) that aligns rigidly the annotation points to a canonical coordinate system (see Fig. 2). Our objective is to find the LV contour with the following decision function:

$$\mathbf{y}^* = \mathbb{E}[\mathbf{y}|I, c = 1, \mathcal{D}] = \int_{\mathbf{y}} \mathbf{y} p(\mathbf{y}|I, c = 1, \mathcal{D}) d\mathbf{y}, \quad (9)$$

where $c = 1$ is a random variable indicating the presence of LV in image I . Eq. 9 can be expanded using

$$p(\mathbf{y}|I, c = 1, \mathcal{D}) = \int_{\vartheta} p(\mathbf{y}|\vartheta, I, c = 1, \mathcal{D}) p(\vartheta|I, c = 1, \mathcal{D}) d\vartheta. \quad (10)$$

The first right-hand side term in (10), representing the non-rigid part of the detection, is defined as follows:

$$p(\mathbf{y}|\vartheta, I, c = 1, \mathcal{D}) = \prod_i p(s_i|\vartheta, I, c = 1, \mathcal{D}), \quad (11)$$




Fig. 3. (a) Lines perpendicular (yellow lines) to annotation points (magenta dots), (b) intensity value profiles (from inside to outside the LV) through perpendicular-lines. Those intensity profiles and respective LV contour location are used to train the regressor of the non-rigid classifier. Figure from [1].




Fig. 4. First images of a subset of the sequences used as training and test sets.

where $p(\mathbf{s}_i|\vartheta, I, c = 1, \mathcal{D})$ represents the probability that the point \mathbf{s}_i is located at the LV contour. Assuming that ψ denotes the parameter vector of the classifier for the non-rigid contour, we compute

$$p(\mathbf{s}_i|\vartheta, I, c = 1, \mathcal{D}) = \int_{\psi} p(\mathbf{s}_i|\vartheta, I, c = 1, \mathcal{D}, \psi)p(\psi|\mathcal{D})d\psi. \quad (12)$$

In practice, we run a maximum a posteriori learning procedure to estimate the model parameters [1], which produces ψ_{MAP} , meaning that in the integral (12) we have $p(\psi|\mathcal{D}) = \delta(\psi - \psi_{\text{MAP}})$, where $\delta(\cdot)$ denotes the Dirac delta function. Also, instead of computing the probability $p(\mathbf{s}_i|\vartheta, I, c = 1, \mathcal{D})$, we train a regressor that indicates the most likely edge location (Fig.3); this roughly means that $p(\mathbf{s}_i|\vartheta, I, c = 1, \mathcal{D}) = \delta(\mathbf{s}_i - \mathbf{s}_i^r(\vartheta, I, c = 1, \mathcal{D}))$, with $\mathbf{s}_i^r(\cdot)$ being the regressor result for the i^{th} contour point.

The second right hand side term in (10) represents the rigid detection, which is denoted as

$$p(\vartheta|I, c = 1, \mathcal{D}) = \mathcal{Z}p(c = 1|\vartheta, I, \mathcal{D})p(\vartheta|I, \mathcal{D}) \quad (13)$$

where \mathcal{Z} is a normalization constant, $p(\vartheta|I, \mathcal{D})$ is a prior on the parameter space, and

$$p(c = 1|\vartheta, I, \mathcal{D}) = \int_{\rho} p(c = 1|\vartheta, I, \mathcal{D}, \rho)p(\rho|\mathcal{D})d\rho, \quad (14)$$

with ρ being the vector of classifier parameters, which is estimated through a maximum a posteriori learning procedure [1], producing ρ_{MAP} . This means that in (14) $p(\rho|\mathcal{D}) = \delta(\rho - \rho_{\text{MAP}})$.

4. SELF-TRAINING AND DETECTION PROCEDURES

In this section, we first introduce the training and test sets, the manual annotation protocol, and then we explain the self-training and detection procedures.

We have two sets of annotated data. The first set contains 400 ultrasound images of the LV, which have been taken from 12 sequences (12 sequences from 12 unhealthy subjects¹ with no overlap), where each sequence contains an average of 34 annotated

¹In the data presented in this paper all the subject suffered from some malfunction in the segments of the endocardium of the LV.

frames. Let us denote this set as \mathcal{D} . This set contains images using the apical two and four-chamber views. The second set, used exclusively for testing, contains two sequences of 80 images, where each sequence has 40 annotated images (2 sequences from 2 unhealthy subjects with no overlap). This set is denoted by \mathcal{T} with sequences A and B . Note that there is no overlap between subjects in sets \mathcal{D} and \mathcal{T} . We worked with one cardiologist, who annotated all images in the database, i.e., sets \mathcal{D} and \mathcal{T} (see Fig. 2 for the an illustration of the annotations). The first image of two sequences from \mathcal{D} and two sequences from \mathcal{T} are shown in Fig. 4.

For the training procedure, consider that the parameters of the discriminative classifier $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$ presented in (1) consists of the parameters ρ and ψ of the rigid (14) and non-rigid (12) classifiers, respectively, as follows: $\boldsymbol{\theta} = [\rho, \psi]$. This classifier is initially trained (*supervised* training) with a subset of \mathcal{D} (in this paper, we consider subsets of sizes $\{2, 6, 10, 20, 50, 100, 200\}$ that are formed by uniformly sampling \mathcal{D}) to maximize $p(\mathcal{Y}|\mathcal{X}, \boldsymbol{\theta})$ [1], which builds $\boldsymbol{\theta}^{(0)} = [\rho_{\text{MAP}}^{(0)}, \psi_{\text{MAP}}^{(0)}]$ (for each subset) in Alg. 1. Given a test sequence in \mathcal{T} the classifier is iteratively trained (*semi-supervised* training) using the detection results from the previous time instant $t - 1$, according to the description of Alg. 1.

For the training of the rigid classifier, we build an image scale space $L(\mathbf{p}, \sigma) = \mathcal{N}(\mathbf{p}, \sigma) * I(\mathbf{p})$, where $\mathcal{N}(\mathbf{p}, \sigma)$ is the Gaussian kernel, $*$ is the convolution operator, $I(\mathbf{p})$ is the input image, σ is the image scale parameter, and \mathbf{p} is the image coordinate. Three separate classifiers (14) are trained; one for each scale $\sigma = \{4, 8, 16\}$ (the values and number of scales were determined based on cross validation of the initial training set at iteration $t = 0$). The positive and negative training sets are defined based on a scale-dependent margin m_σ that increases by a factor of two after each octave. Positives for $L(\mathbf{p}, \sigma)$ are randomly generated *inside* the range $[\vartheta - m_\sigma/2, \vartheta + m_\sigma/2]$, and negatives are randomly generated *outside* the range $[\vartheta - m_\sigma, \vartheta + m_\sigma]$, where ϑ is the parameter vector representing the rigid transformation of the LV annotation. The non-rigid regressor (12) is trained at $\sigma = 4$, where each training sample is a line of 41 pixels of length extracted perpendicularly from the LV contour points (see Fig. 3) and the label to learn is the pixel index in $\{1, \dots, 41\}$ that is closest to the LV contour. A cross-validation procedure using 20% of the initial training set for validation is used to estimate the following parameters: 1) number of nodes per layer of regressor network; 2) number of nodes per layer of the classifier networks; and 3) the prior distribution $p(\vartheta|I, \mathcal{D})$ used in (13).

The detection procedure consists of running the rigid classifier at scale $\sigma = 16$ on the K_{coarse} initial hypotheses [1] (here, $K_{\text{coarse}} = 1000$), by sampling the random distribution $p(\vartheta|I, \mathcal{D})$ from (13). From this detection, cluster the hypotheses (using k-means algorithm) and select the top K_{fine} clusters (here, $K_{\text{fine}} = 10$) in terms of the best hypothesis within each cluster [1]. Then run the rigid classifier at scale $\sigma = 8$ on these hypotheses and repeat the procedure for scale $\sigma = 4$. Finally, run the model represented by (10) over the final top K_{fine} hypotheses. Note that we substitute the integral in (9) for an average over the top K_{fine} hypotheses weighted by $p(\mathbf{y}|I, c = 1, \mathcal{D})$. The final segmentation contour points are then projected to the principal component analysis (PCA) space built with the respective subset of the training set \mathcal{D} . The PCA space transforms the 41-dimensional vector (representing the contour) to a 5-dimensional vector, which is back projected onto the original contour space, producing a less noisy final contour [1].

5. EXPERIMENTS

In this section, the importance in the number of images used to estimate $\boldsymbol{\theta}^{(0)}$ in Alg. 1. We also compare quantitatively the performance of the supervised and semi-supervised methods. Fur-




Fig. 5. Examples of the detection improvement provided by the semi-supervised (right column) learning compared to the supervised (left column) model trained with 50 images.

thermore, we compare our results to state-of-the-art LV detectors recently proposed by Carneiro et al. [1], by Comaniciu et al. [2,3] and by Nascimento et al. [8]. The performance of the detectors is assessed by comparing the contour estimates with manual reference contours (see Sec. 4). To accomplish this, we use four error metrics used in the literature: (i) average metric (AV) (ii) Hausdorff metric (HDF), (iii) Hammoude distance (HMD) and (iv) mean absolute distance (MAD) (see [2,8] for a comprehensive description of these error metrics).

Now we show how the semi-supervised training method improves the performance of the system initially trained with small training sets (this initial system is labeled 'Supervised'). We also compare the results with the performance of the methods proposed in [1], [3] and [8]. For this experiment, we build three different training sets of sizes in S and show the results using mean and standard deviation for each error measure (Fig. 6). Compared to the supervised training, note that the proposed semi-supervised learning reduces both the mean and the standard deviation for most of the error measures. In general, the semi-supervised approach starts producing state-of-the-art results with initial training sets containing 50 images, but notice that for sequence $T(A)$ the system shows competitive results with initial training sets containing only 6 images. Figure 5 displays two cases showing the improvement provided by semi-supervised learning.

6. DISCUSSION AND CONCLUSIONS

In this paper, we presented a novel semi-supervised self-training methodology applied to the segmentation of the left ventricle of the heart from ultrasound data. The novelties reside in the formulation of the self-training algorithm that keeps adding training images as frames of a new test sequence are presented to the system. This means that the initial set of annotated and unannotated training images is not fixed, which is a common assumption adopted by semi-supervised learning approaches. For this reason, the selection criterion to add unannotated images to the training set becomes a critical aspect of the algorithm, and we provide an empirical study on the selection of such criterion. We also derived the formulation of our algorithm. The results show that it is possible to have state-of-the-art results with training sets containing 50 annotated training images. We plan to study better selection criterion methods [9] to improve even more the results presented in this paper.

7. REFERENCES

[1] G. Carneiro, J. Nascimento, and A. Freitas, "Robust left ventricle segmentation from ultrasound data using deep neural networks and efficient search methods," in *ISBI*, 2010.
 [2] D. Comaniciu, X. Zhou, and S. Krishnan, "Robust real-time myocardial border tracking for echocardiography: An information fusion approach," *IEEE Trans. Med. Imag.*, vol. 23, no. 7, pp. 849–860, 2004.

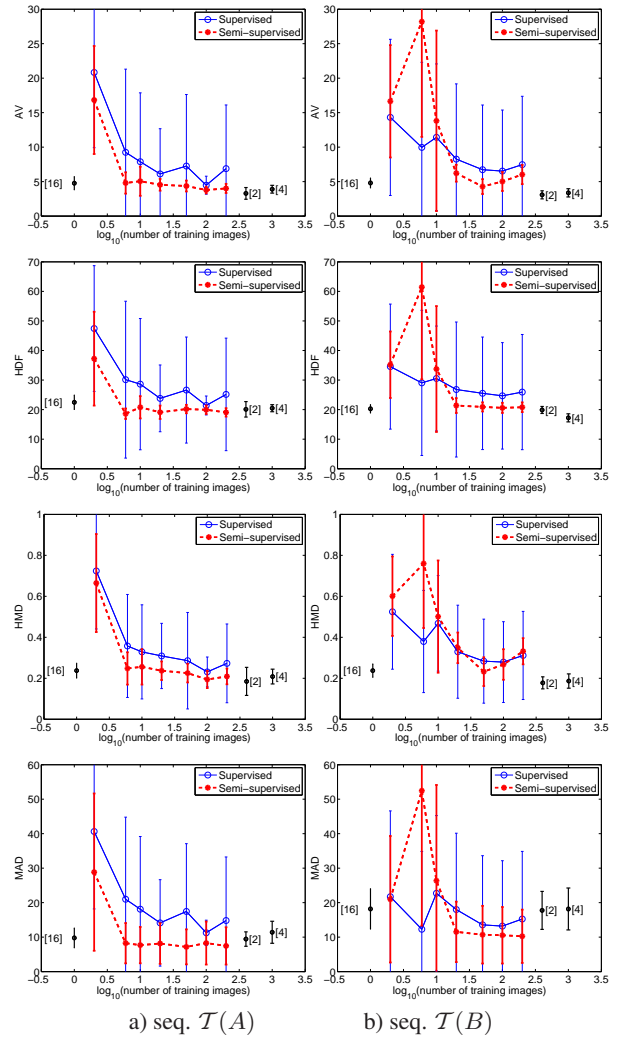


Fig. 6. Comparison of the performance of the proposed semi-supervised method and the supervised approach using the four error measures (each row represents one error measure, and each column denotes a different test sequence). We also show the detection results on the same test sets of the supervised training methods [1] and [3] and the unsupervised model-based method [8].

[3] B. Georgescu, X. S. Zhou, D. Comaniciu, and A. Gupta, "Database-guided segmentation of anatomical structures with complex appearance," in *Conf. Computer Vision and Pattern Rec. (CVPR)*, 2005.
 [4] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu, "Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features," *IEEE Trans. Med. Imaging*, vol. 27, no. 11, pp. 1668–1681, 2008.
 [5] Xiaojin Zhu, "Semi-supervised learning literature survey," Tech. Rep. 1530, Computer Sciences, University of Wisconsin-Madison, 2005.
 [6] R. Neal and G. Hinton, "A view of the em algorithm that justifies incremental, sparse, and other variants," in *Learning in Graphical Models*. 1998, pp. 355–368, Kluwer Academic Publishers.
 [7] G. Hinton and R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
 [8] J. C. Nascimento and J. S. Marques, "Robust shape tracking with multiple models in ultrasound images," *IEEE Trans. Imag. Proc.*, vol. 17, no. 3, pp. 392–406, 2008.
 [9] C. Rosenberg, M. Hebert, and H. Schneiderman, "Semi-supervised selftraining of object detection models," in *Seventh IEEE Workshop on Applications of Computer Vision*, 2005.