Accelerating the Optimization of a Segmentation Ensemble using Image Pyramids

János Tóth, Tibor Péter Kapusi, Balázs Harangi, Henrietta Tomán, and András Hajdu
Faculty of Informatics, University of Debrecen
4002 Debrecen, POB 400, Hungary
e-mail: {toth.janos, kapusi.tibor, harangi.balazs, toman.henrietta, hajdu.andras}@inf.unideb.hu

***Abstract*—In this paper, we propose an image pyramid-based noisy energy function evaluation method for the local search technique simulated annealing. The method is primarily designed for the optimization of image segmentation algorithms, and it maintains solution quality with significantly reduced time requirement. The strategy to select the proper image pyramid levels during the search is theoretically determined via adapting results regarding evaluation in simulated annealing based on imprecise measurements. As a demonstrative application, we perform parameter-optimization of a segmentation ensemble dedicated to the extraction of bone structures from CT images.**

***Index Terms*—ensemble-based system, segmentation, noisy evaluation, image pyramid**

I. Introduction

Optimization of a system having numerous free parameters regarding a complex energy function is a very challenging task. Even if we can transform the problem to a finite discrete one via e.g. quantizing and limiting the domains of the parameters, the number of possible parameter settings grows exponentially with the number of the parameters. Accordingly, an exhaustive search for the optimal parameter setting soon becomes completely impractical. Moreover, the more efficient stochastic search approaches may still require too much time, especially if the evaluation of the energy function itself is expensive. In this paper, we address this issue by suggesting a special type of noisy evaluation [1], [2] of the energy function of simulated annealing (SA) [3], [4], which is a popular and widely applied stochastic search strategy.

As a specific application, we consider an ensemble of segmentation algorithms dedicated to the extraction of bone structures from computer tomography (CT) images. The outputs of the individual segmentation algorithms are binary images containing the candidate bone regions, which are aggregated by majority voting. These algorithms have numerous adjustable parameters, whose optimal settings may differ, when a given algorithm is considered as an individual approach or a member of an ensemble. The evaluation of the ensemble performance for a given parameter setting is done by comparing the aggregated segmentation results of the ensemble for a training set with the corresponding manually annotated ground truth. The comparison is carried out using the statistic intersection over union (IoU) [5], which is commonly considered for object detection. Thus, to compute the energy function during the optimization process, an evaluation over the whole training set should be performed at each search step. A possible speed up during the stochastic search is to apply noisy evaluation to determine the ensemble performance, which means that the energy function is just approximated instead of determined precisely. Naturally, to maintain the convergence characteristics of the search method the noise should be controlled. In the case of SA, it requires that the noise should be normally distributed with mean 0 and its variance should tend to 0 as the search advances [1].

In [6], we have already successfully presented how such a noisy evaluation can be provided by selecting only subsets of the training set with proper cardinalities instead of the whole training set during the search. Now, as a novel contribution, we consider an image pyramid representation of the training set, where evaluating on lower resolution levels results in noisy determination of the energy. Naturally, the lower the resolution is, the larger the noise can be, since the segmentation results using lower resolution images can be less precise. To meet the theoretical requirements, we introduce a strategy to determine the maximal allowed noise level in each iteration to control the search. We will show that our method successfully reduces the time requirement of the search while preserving solution quality.

The rest of the paper is organized as follows. In section II we give an overview of the individual segmentation algorithms together with their adjustable parameters and the aggregation strategy to construct an ensemble from them. The image pyramid based evaluation method and the strategy to select the appropriate scaling levels to guarantee the theoretical requirements for the noise allowed in the evaluation during the search is given in section III. Our experimental setup and results regarding the performance of the proposed evaluation method for the optimization of the segmentation ensemble will be provided in section IV. Finally, some conclusions are drawn in section V.

II. An Ensemble for Automatic Bone Segmentation in CT Images

In this section, we present an ensemble that performs automatic bone segmentation in CT images. We describe its member algorithms and the aggregation method used to

This work was supported in part by the Janos Bolyai Research Scholarship of the Hungarian Academy of Sciences and the projects EFOP-3.6.2-16-2017-00015 and GINOP-2.2.1-15-2017-00055 supported by the European Union, co-financed by the European Social Fund.

generate the output of the ensemble based on the individual output of the members.

A. Basic notations

Let $A_1, A_2, \ldots, A_N$ denote the individual segmentation algorithms of the ensemble, and $C_1, C_2, \ldots, C_N$ be their corresponding output. Let $C_i(x, y)$ denote the pixels of the output of A_i, where $i = 1, \ldots, N$. Furthermore, let $p_1, p_2, \ldots, p_n$ be the parameters of the ensemble, where n is the total number of parameters.

B. Member algorithms

Our automatic bone segmentation ensemble consists of five algorithms, each of which has many parameters. However, to gain a problem that is computationally reasonable, we selected only those parameters for the later optimization that significantly influence the ensemble output.

1) Algorithm A_1: The algorithm A_1 uses distance regularized level set evolution (DRLSE) [7] with thresholding initialization and an edge-based active contour model. A_1 has the following parameters: the coefficient of the weighted area term, the width of the Dirac δ function, the coefficient of the weighted length term and a time-step parameter, which affects the coefficient of the distance regularization term.

Among these parameters, the width of the δ function (p_1) is the most relevant regarding the accuracy of the output.

2) Algorithm A_2: This algorithm is based on the dual threshold technique described in [8] for extracting the periosteal and endosteal surfaces of the bones in two steps. We have implemented only the first step of the method for extracting the bone surface from CT images. The algorithm A_2, that applies thresholding and morphological operations, has the following parameters: the number of thresholding levels, the size of the median filter, and the parameters of the morphological structuring elements. In the case of A_2, we chose the number of thresholding levels (p_2) for the optimization.

3) Algorithm A_3: The algorithm A_3 uses fuzzy C-means clustering [9]. In the last step, Hounsfield-unit based thresholding of the input image is performed, and the clustering result having the least symmetric difference compared to the Hounsfield output is selected.

The range for thresholding has been selected to be 500 to 900 HU [10]. A_3 has the following parameters: the number of clusters, the exponent for the fuzzy partition matrix, the iteration number and the improvement value of the objective function. Among these parameters, the number of clusters (p_3) and the exponent (p_4) have the largest influence on the output.

4) Algorithm A_4: The algorithm A_4 [11] performs histogram matching, morphological operations, and finally active contour segmentation using the method developed by Chan and Vese [12]. A_4 has the following parameters: the number of thresholding levels, the parameters of the morphological structuring elements, the weight of the smoothing term, and the number of iterations for the active contour segmentation. Among these parameters, the number of thresholding levels (p_5) has the most significant influence on the output.

5) Algorithm A_5: This algorithm is a variant of the region growing method [13] with multiple seed points. It compares iteratively the intensity of each unallocated neighboring pixel to the mean of the already segmented region until the difference of these values becomes larger than a threshold. Initial seed points are selected using the histogram of the input image, and the similarity threshold is automatically estimated using the variance of the input image. A_5 has two parameters: the number of initial seed points, and a correction factor of the similarity threshold, of which we chose the latter (p_6) for the optimization.

In Table I, we summarize the adjustable parameters of the ensemble members.

TABLE I
ADJUSTABLE PARAMETERS OF THE ENSEMBLE MEMBERS.

Alg.	Parameter description	Range
A_1	width of δ function	$p_1 \in \{0.01, 0.21, \ldots, 2.01\}$
A_2	thresholding levels	$p_2 \in \{2, 3, \ldots, 5\}$
A_3	number of clusters	$p_3 \in \{2, 3, \ldots, 7\}$
A_3	exponent	$p_4 \in \{1.01, 1.21, \ldots, 3.41\}$
A_4	thresholding levels	$p_5 \in \{2, 3, \ldots, 5\}$
A_5	correction factor	$p_6 \in \{0.8, 0.9, \ldots, 1.2\}$

C. Aggregation method

As the aggregation method to obtain the output of the ensemble, we chose classic majority voting. That is, the pixel values of the ensemble output C is determined as

$$C(x,y) = \begin{cases} 1, & if \sum_{i=1}^{N} C_i(x,y) \geq \left[\frac{N}{2}\right] \\ 0, & otherwise \end{cases}, \qquad (1)$$

where N = 5 is the number of member algorithms.

See Fig. 2 for examples of the output of the individual algorithms and the ensemble for a CT image.

III. ACCELERATED PARAMETER OPTIMIZATION USING IMAGE PYRAMID-BASED NOISY EVALUATION

In this section, we present a novel image pyramid-based noisy energy function evaluation method for the parameter optimization of image segmentation ensembles that reduces the time requirement of SA while preserves its convergence properties.

A. Simulated Annealing in Presence of Noise

Simulated annealing is a local search strategy [3], [14] that is widely applied to address difficult combinatorial optimization problems. The main feature of SA is that it can escape from local optima by allowing moves that deteriorate the energy function value with a probability depending on a control parameter and the energy function difference of the candidate and the current solution. SA assumes that the energy of a solution can be determined exactly; however, the evaluation of a solution is often subject to noise in real-life problems.

Considering discrete search spaces and assuming that in the k-th ($k \in \mathbb{N}$) iteration the noise is normally distributed with mean 0 and variance $\left(\sigma^{(k)}\right)^2 > 0$, Gelfand and Mitter proved [1] that SA using noisy evaluation exhibits the same convergence properties as using exact energy values, if the standard deviation $\sigma^{(k)}$ of the noise is dominated by the temperature $T^{(k)}$ in the k-th iteration for each k, that is, when

$$\sigma^{(k)} = o\left(T^{(k)}\right). \tag{2}$$

Next, we will describe how the criterion defined by (2) can be exploited to accelerate the parameter optimization of an image segmentation ensemble using image pyramids.

B. Image Pyramid-based Noisy Evaluation in SA

In the case of ensembles, the more parameters they have, the larger the training set required is to find their optimal parameter setting and to avoid overfitting. However, considering large training sets, evaluating the mean performance of the ensemble with a parameter setting can be computationally expensive even using simple cost functions. Therefore, we propose to estimate the energy function value E of SA at different downscaling levels of the original ground truth and the aggregated output of the ensemble during the SA process.

1) Nearest neighbor image pyramid: We refer to a collection of $L \in \mathbb{N}$ hierarchically downscaled versions of an image as an L-level image pyramid, in which the higher the level l ($l \in 0, 1, \dots L-1$), the smaller the image resolution is. See Fig. 1 for a visual explanation of the construction.

The most common method to construct image pyramids is the Gaussian pyramid [15], where levels of the pyramid are built by convolving the original image with a Gaussian-like averaging filter followed by a subsampling step. However, in our case the images are binary masks, therefore to preserve sharp boundaries, we use the nearest neighbor method for simple subsampling. That is, the pixel values of a level are defined to match the original pixel whose center is the nearest to the sample position.

2) The effects of downscaling: Assuming that the cost of calculating E is proportional to the resolution of the input images, the calculation of the energy function estimate $\widehat{E}_l$ over the l-th level version of the input images (with an associated scaling factor γ_l) has $1/\gamma_l^2$ times lower cost than the calculation of E; however, using $\widehat{E}_l$ introduces noise in the evaluation. The noise d_l originating from using the l-th level version of the input images can be determined as

$$d_l = \widehat{E}_l - E. \tag{3}$$

This noise may cause SA to consider an inferior state as superior because of the imprecise evaluation of the energy function. That is, the stronger the noise, the more random the search. For this reason, a suitable strategy to control the noise is required.

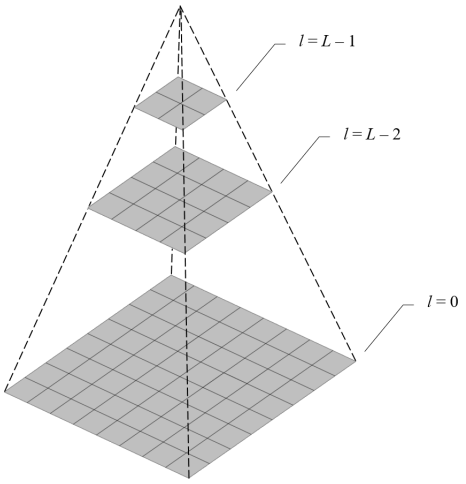

Fig. 1. Image pyramid

3) Downscaling level selection strategy: To ensure the convergence of SA, we have to apply a strategy to select the appropriate downscaling level l in each iteration k to control the standard deviation of the noise $\sigma_d^{(k)}$ according to the temperature $T^{(k)}$, as described by (2).

First, we have to determine the maximum allowed value of $\sigma_d^{(k)}$ for a given temperature $T^{(k)}$. Using (2), we get that

$$\lim_{k\to\infty} \frac{\sigma_d^{(k)}}{T^{(k)}} = 0 \tag{4}$$

must hold.

To maintain the limit in (4), the sequence $\left\{\sigma_d^{(k)}\right\}$ has to be decreasing such that $\lim_{k\to\infty} \sigma_d^{(k)} = 0$ and $\sigma_d^{(k)} < T^{(k)}$ for each $k \in \mathbb{N}$.

Based on the above conditions, a sufficiently simple general form of $\sigma_d^{(k)}$ that maximizes its value considering a given $T^{(k)}$ can be given as:

$$\sigma_d^{(k)} = T^{(k)}(1-\epsilon)^k \text{ with } 0 < \epsilon < 1. \tag{5}$$

The next step to construct our strategy is to determine the dependence of $\sigma_d^{(k)}$ on the downscaling level l.

Using a downscaled version of the images of the training set during the evaluation may result in significantly different noise for different energy functions. In some cases, it can be straightforward to determine the theoretical maximum value of $\sigma_d^{(k)}$, while for more complex energy functions it becomes a difficult problem. However, in the case of natural images the empirical standard deviation of the noise $\sigma_{d,l}$ for a level l is expected to be much lower than the theoretical maximum. Therefore, we propose to estimate $\sigma_d^{(k)}$ for each level of the image pyramid by measuring the value of $\sigma_{d,l}$ on the ground truth used for the evaluation.

Having $\sigma_{d,l}$ measured for each level l of the image pyramid, we can determine the highest level l (i.e. the lowest resolution) where $\sigma_{d,l}$ is less than or equal to the maximum allowed $\sigma_d^{(k)}$ for each temperature level $T^{(k)}$.

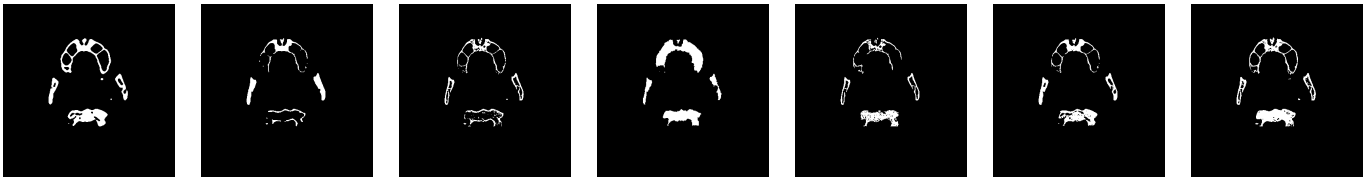

(a) (b) (c) (d) (e) (f) (g)

Fig. 2. Examples of the algorithm outputs, the aggregated output of the ensemble, and the corresponding ground truth: (a) output of A_1, (b) output of A_2, (c) output of A_3, (d) output of A_4, (e) output of A_5, (f) ensemble output, (g) ground truth.

TABLE II
RESULTS FOR THE DATASET.

		p_1	p_2	p_3	p_4	p_5	p_6	$\overline{SE}$	$\overline{SP}$	$\overline{MCC}$	$\overline{ACC}$	$\overline{IoU}$	t (seconds)
Individual evaluation	test	1.61	4	3	1.61	3	1.2	0.8519	0.9984	0.8827	0.9920	0.8238	-
Exhaustive search	training	1.81	4	3	1.21	3	1.2	0.9090	0.9970	0.9200	0.9932	0.8597	71343.7 (19.818 hours)
	test							0.8588	0.9980	0.8826	0.9919	0.8235	-
SA	training	0.21	4	3	1.21	3	1.2	0.9085	0.9970	0.9197	0.9931	0.8591	137.3
	test							0.8583	0.9980	0.8823	0.9919	0.8230	-
SA with noisy evaluation	training	0.21	4	3	1.41	3	1.2	0.9021	0.9975	0.9201	0.9933	0.8589	48.9
	test							0.8509	0.9984	0.8821	0.9919	0.8226	-

Next, we present our experimental setup and our results to evaluate the performance of the proposed method.

IV. EXPERIMENTAL RESULTS

In this section, we describe the methodology used to assess the performance of the proposed image pyramid-based noisy evaluation method, and present our quantitative results, with highlighting how the downscaling of the images affects the optimization process.

A. SA Design Choices

1) Energy function: To evaluate the proposed noisy evaluation method, we have performed parameter optimization of the ensemble presented in section II, with the aim of efficiently finding the parameter setting that maximizes the segmentation performance of the ensemble in terms of average intersection over union $\overline{IoU}_l$ [5], which is defined as the number of common pixels of the ensemble output $C_{m,l}$ and the corresponding ground truth $G_{m,l}$ for a given level l of the image pyramid ($l = 0, 1 \ldots, L-1, \quad m = 1, 2, \ldots, M$) over the number of pixels in either of the two:

$$\overline{IoU}_l = \sum_{i=1}^{M} \frac{|C_{m,l} \cap G_{m,l}|}{|C_{m,l} \cup G_{m,l}|}. \tag{6}$$

To obtain a minimization problem, we define the energy function as:

$$E = 1 - \frac{\overline{IoU}_0}{M}, \tag{7}$$

and the energy function estimate for the level l as:

$$\widehat{E}_l = 1 - \frac{\overline{IoU}_l}{M} \tag{8}$$

where L is the number of levels in the image pyramid and M is the number of images in the training set.

2) Cooling schedule: For the implementation of the search, we chose the exponential cooling schedule via

$$T^{(k)} = T^{(0)}\, \alpha^k \text{ with } 0 \leq \alpha \leq 1. \tag{9}$$

We set the initial temperature $T^{(0)} = 1$ and the base $\alpha = 0.985$. As the stopping criterion, we chose to have fixed number of iterations, with $k_{max} = 500$.

3) Dataset: Our dataset consists of 300 private cross-sectional CT images taken of the head of one patient and the corresponding manually annotated ground truth masks. Of course, the proposed method is not specific to any anatomic location, it could be adapted for other CT images as well. The dataset was randomly divided into two parts: a training set with 200, and a test set with 100 images. The images have the resolution of 512×512 pixels. The image pyramid representation of the dataset was constructed using $L = 16$ levels, with corresponding scaling factor γ_l, defined as

$$\gamma_l = 1 - l/L, \quad l = 0, 1, \ldots, L-1. \tag{10}$$

4) Realization of the noisy evaluation: Using the setup and dataset described above, we have measured the standard deviation of the noise $\sigma_{d,l}$ for each level l over the ground truth of the training set, and for comparison over the aggregated output of the ensemble. For each level l, we used the followings method:

- Downscale the ground truth and ensemble output images with scaling factor γ_l to construct the corresponding level of the image pyramids.

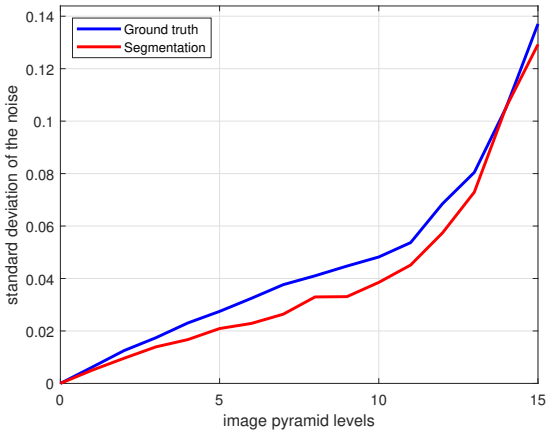

Fig. 3. Measured standard deviation of the noise for the training set

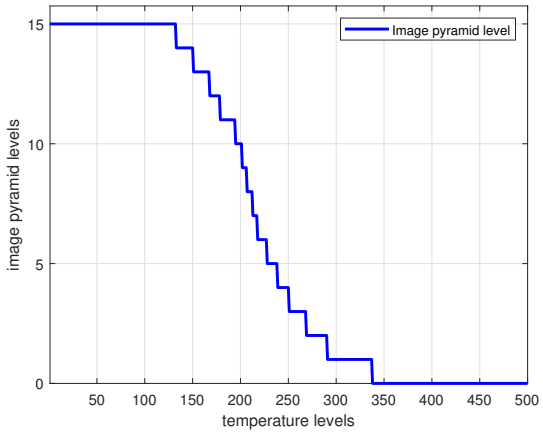

Fig. 5. Required image size for a given temperature level during the search

- Upscale the images to the original size (with scaling factor $1/\gamma_l$).
- Determine d_l as defined in (3). For this, compute E and its estimate $\widehat{E}_l$ for the level l as described in IV-A1 using the scaled versions of the m-th image (for each m) as $C_{m,l}$ and the original images as $G_{m,l}$.
- Calculate $\sigma_{d,l}$.

It can be observed in Fig. 3 that the standard deviation of the noise exhibits a similar trend in the case of both the ground truth and the corresponding ensemble output.

Using (9) the maximum allowed value of $\sigma_d^{(k)}$ can be calculated as

$$\sigma_d^{(k)} = T^{(0)}\,\alpha^k(1-\epsilon)^k \text{ with } 0 \leq \alpha \leq 1, \text{ and } 0 < \epsilon < 1. \tag{11}$$

Using (11), we were able to determine the required level l for each iteration (temperature level) k. The maximal allowed standard deviation of the noise $\sigma_d^{(k)}$, the fitted noise $\sigma_{d,l}$, and the corresponding downscaling levels l are shown in Fig. 4 and in Fig. 5, respectively.

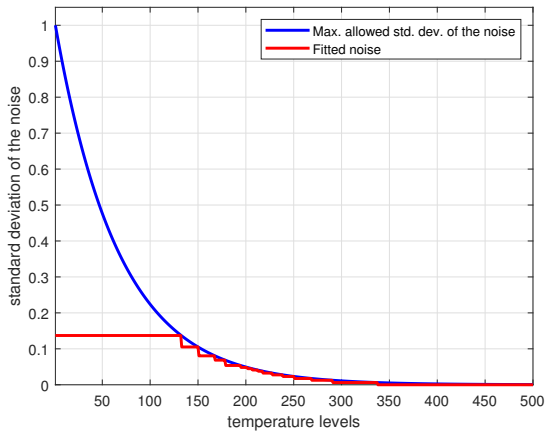

Fig. 4. Maximal fitted standard deviation of the noise

5) Quantitative results: In Table II, we list the individually optimal parameter values, and the optimal parameter values found at ensemble-level using exhaustive search, SA, and SA with the proposed noisy evaluation method, along with the corresponding average sensitivity ($\overline{SE}$), specificity ($\overline{SP}$), Matthews correlation coefficient ($\overline{MCC}$), accuracy ($\overline{ACC}$), and IoU ($\overline{IoU}$) values of the ensembles using the training set, as well as the running times (t) in seconds required for the parameter optimization. The running time for the individually optimal parameter setup and for the test set are omitted since in these cases only evaluation is performed.

In order to assess the stability of the method, we performed 300 tests using both standard SA and SA with the proposed noisy evaluation method. In Table III we include the average figures of 300-300 tests. As it can be seen, the algorithm exhibits a solid behavior for the stability of the search with small differences in the average energy function values; however, significant improvement in the time requirement is achieved.

TABLE III
PERFORMANCE COMPARISON BASED ON THREE HUNDRED RUNS.

	$\overline{IoU}$	t (seconds)
SA	0.8248	119.5
SA with noisy evaluation	0.8200	39.6
Difference	-0.0048 (-0.58%)	-79.9 (-66.86%)

6) Implementation and hardware details: The algorithms were implemented in Matlab. All detector outputs and ground truth images represented as image pyramids are stored in memory during the optimization process to reduce the time required to find a solution. The reported running times exclude the time required for loading the ground truth images and the algorithm output that were computed offline, generating the image pyramids, and other overhead. Results for the dataset were acquired using a computer equipped with a 4-core 8-thread Intel Xeon W-2123 processor and 16 GB DDR4 RAM.

V. Conclusions

In this paper, we have proposed an image pyramid-based noisy energy function evaluation method for the local search technique SA. Considering an image segmentation ensemble designed to extract bone structures from CT images, we showed that using the proposed method it is possible to reach solutions with the same quality as using standard simulated annealing, but with significantly reduced time requirement.

References

[1] S. B. Gelfand and S. K. Mitter, "Simulated annealing with noisy or imprecise energy measurements," *Journal of Optimization Theory and Applications*, vol. 62, no. 1, pp. 49–62, 1989. [Online]. Available: https://doi.org/10.1007/BF00939629

[2] P. C. Schuur, "Classification of acceptance criteria for the simulated annealing algorithm," *Mathematics of Operations Research*, vol. 22, no. 2, pp. 266–275, 1997. [Online]. Available: http://www.jstor.org/stable/3690264

[3] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983. [Online]. Available: https://doi.org/10.1126/science.220.4598.671

[4] S. Ledesma, G. Avia, and R. Sanchez, "Practical considerations for simulated annealing implementation," in *Simulated Annealing*, C. M. Tan, Ed. Vienna, Austria: InTech, 2008, ch. 20, pp. 401–420. [Online]. Available: http://doi.org/10.5772/5560

[5] W. R. Crum, O. Camara, and D. L. G. Hill, "Generalized overlap measures for evaluation and validation in medical image analysis," *IEEE Transactions on Medical Imaging*, vol. 25, no. 11, pp. 1451–1461, Nov 2006. [Online]. Available: https://doi.org/10.1109/TMI.2006.880587

[6] J. Tóth, L. Szakács, and A. Hajdu, "Finding the optimal parameter setting for an ensemble-based lesion detector," in *2014 IEEE International Conference on Image Processing (ICIP)*, Oct 2014, pp. 3532–3536. [Online]. Available: https://doi.org/10.1109/ICIP.2014.7025717

[7] C. Li, C. Xu, C. Gui, and M. D. Fox, "Distance regularized level set evolution and its application to image segmentation," *Trans. Img. Proc.*, vol. 19, no. 12, pp. 3243–3254, Dec. 2010. [Online]. Available: http://doi.org/10.1109/TIP.2010.2069690

[8] H. R Buie, G. Campbell, R. Joshua Klinck, J. A MacNeil, and S. K Boyd, "Automatic segmentation of cortical and trabecular compartments based on a dual threshold technique for in vivo micro-CT bone analysis," *Bone*, vol. 41, pp. 505–15, 11 2007. [Online]. Available: https://doi.org/10.1016/j.bone.2007.07.007

[9] H. Zhou, G. Schaefer, and C. Shi, *Fuzzy C-Means Techniques for Medical Image Segmentation*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 257–271. [Online]. Available: https://doi.org/10.1007/978-3-540-89968-6_13

[10] D. Lim Fat, J. Kennedy, R. Galvin, F. O'Brien, F. Mc Grath, and H. Mullett, "The hounsfield value for cortical bone geometry in the proximal humerus—an in vitro study," *Skeletal Radiology*, vol. 41, no. 5, pp. 557–568, May 2012. [Online]. Available: https://doi.org/10.1007/s00256-011-1255-7

[11] M. Kardell, M. Magnusson, M. Sandborg, G. Alm Carlsson, J. Jeuthe, and A. Malusek, "Automatic segmentation of pelvis for brachytherapy of prostate," *Radiation Protection Dosimetry*, vol. 169, no. 1-4, pp. 398–404, 06 2016. [Online]. Available: https://doi.org/10.1093/rpd/ncv461

[12] T. F. Chan and L. A. Vese, "Active contours without edges," *Trans. Img. Proc.*, vol. 10, no. 2, pp. 266–277, Feb. 2001. [Online]. Available: https://doi.org/10.1109/83.902291

[13] R. M. Haralick and L. G. Shapiro, "Image segmentation techniques," *Computer Vision, Graphics, and Image Processing*, vol. 29, no. 1, pp. 100 – 132, 1985.

[14] V. Černý, "Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm," *Journal of Optimization Theory and Applications*, vol. 45, no. 1, pp. 41–51, 1985. [Online]. Available: http://doi.org/10.1007/BF00940812

[15] P. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Transactions on Communications*, vol. 31, no. 4, pp. 532–540, April 1983. [Online]. Available: https://doi.org/10.1109/TCOM.1983.1095851