Joint graph cut and relative fuzzy connectedness image segmentation algorithm ☆

Krzysztof Chris Ciesielski [a,b,*], Paulo A.V. Miranda [c], Alexandre X. Falcão [d], Jayaram K. Udupa [b]

[a] Department of Mathematics, West Virginia University, Morgantown, WV 26506-6310, United States
[b] Department of Radiology, MIPG, University of Pennsylvania, Blockley Hall, 4th Floor, 423 Guardian Drive, Philadelphia, PA 19104-6021, United States
[c] Department of Computer Science, IME, University of São Paulo (USP), São Paulo, SP, Brazil
[d] Institute of Computing, University of Campinas, Campinas, SP, Brazil

A B S T R A C T

We introduce an image segmentation algorithm, called GC_{sum}^{max}, which combines, in novel manner, the strengths of two popular algorithms: Relative Fuzzy Connectedness (RFC) and (standard) Graph Cut (GC). We show, both theoretically and experimentally, that GC_{sum}^{max} preserves robustness of RFC with respect to the seed choice (thus, avoiding "shrinking problem" of GC), while keeping GC's stronger control over the problem of "leaking though poorly defined boundary segments." The analysis of GC_{sum}^{max} is greatly facilitated by our recent theoretical results that RFC can be described within the framework of Generalized GC (GGC) segmentation algorithms. In our implementation of GC_{sum}^{max} we use, as a subroutine, a version of RFC algorithm (based on Image Forest Transform) that runs (provably) in linear time with respect to the image size. This results in GC_{sum}^{max} running in a time close to linear. Experimental comparison of GC_{sum}^{max} to GC, an iterative version of RFC (IRFC), and power watershed (PW), based on a variety medical and non-medical images, indicates superior accuracy performance of GC_{sum}^{max} over these other methods, resulting in a rank ordering of GC_{sum}^{max} > PW $\sim$ IRFC > GC.

1. Introduction

The image segmentation field has a rich literature dating back to the 1960s. The algorithms for image segmentation can be categorized into three groups: purely image-based (pI), shape model-based (SM), and hybrid. pI methods (Liang et al., 2006; Malladi et al., 1995; Udupa and Samarasekera, 1996; Singaraju et al., 2010; Pham, 2001; Boykov et al., 2001; Kolmogorov and Zabih, 2004) focus on delineating objects based entirely on the information about the object that can be harnessed from the given image. SM approaches (Cootes et al., 1999; Sandor and Leahy, 1997; Pizer et al., 2003) bring in information about the object family in terms of its appearance variation in the form of statistical/fuzzy texture and/or shape models to bear on the segmentation problem. Hybrid approaches (Miranda et al., 2009; Liu and Udupa, 2009; Imielinska et al., 2001; Angelini-Casadevall et al., 2002; Lim, 2006) are recent; they combine synergistically the pI and SM approaches in an attempt to overcome the weaknesses of the individual approaches.

The major frameworks existing under the pI approaches include: level sets (LS), active boundaries, clustering, Markov Random Field, Graph Cut (GC), fuzzy connectedness (FC), and watershed (WS).

All these classes of approaches have their place in the segmentation armamentarium. The focus of this paper is on the pI approaches. Since they are based entirely on the information available in the given image, and since top-rated algorithms in this group harness this information with equal effectiveness, there must exist similarity or even equivalence among such algorithms. This observation prompted researchers to study the possibility of explaining such algorithms in a common framework (Ciesielski and Udupa, 2011a; Aubert and Blanc-Féraud, 1999; Miranda and Falcão, 2009). In the same spirit, the popular GC framework has been generalized recently to, what we refer to as, *Generalized GC* (GGC). This framework was proposed by the authors in (Ciesielski and Udupa, 2011b; Ciesielski et al., 2011, 2012b), and studied in a slightly different form in (Couprie et al., 2011), to describe GC, FC, and WS algorithms in a unified manner. A byproduct of such a unification effort was a deeper understanding of the strengths and weaknesses of the individual algorithms, which can lead to new methods with improved performance, as we will demonstrate in this paper.

Despite the success, attested in several applications, and the popularity of approaches such as watersheds (Miranda and Falcão, 2009; Couprie et al., 2011; Sinop and Grady, 2007), iterative relative fuzzy connectedness (Ciesielski et al., 2007), optimum-path

forest (Miranda and Falcão, 2009; Falcão et al., 2004; Falcão and Bergo, 2004), and shortest paths (Bai and Sapiro, 2007; Protiere and Sapiro, 2007), the lack of regularity constraints often makes these approaches inappropriate (or at least not the best choice) for some application domains. For instance, whenever the object is expected to present simple and regular shapes, the presence of poorly defined boundary segments can cause leaking problems in the above methods, resulting often in not only wrong but also irregular (jagged) boundaries.

In order to amend this problem, some authors consider the usage of regularization energies (e.g., internal forces) intrinsic to their formulations (Boykov et al., 2001; Malladi et al., 1995; Shi and Malik, 2000; Grady, 2006), while others enforce smoothness by post-processing (Malmberg et al., 2010, 2011; Malmberg, 2011; Falcão et al., 2002). The main drawback of the second group of methods is that it may be too late to fix a result, when it is already too far from the goal (e.g., significant leakage). However, such methods may be beneficial in applications where only some specific parts of the object are supposed to be smooth. In this case, we can easily constrain the post-processing to those areas, while global intrinsic regularization energies will affect other object parts which are not necessarily intended to be smooth. On the other hand, the methods with intrinsic regularization are less prone to leakage. However this usually comes at a price, such as the shrinking problem (GC), local optima issues (LS, snakes), poor robustness to seed position (GC, random walker), besides the higher computational cost.

This work falls within the GGC framework described in (Ciesielski et al., 2011, 2012b), wherein we identified and justified some of the strong and weak properties of GC and FC, both theoretically and empirically, in a comparative manner. The most crucial among these were robustness of segmentation with respect to the selection of seed points (FC better than GC), boundary smoothness (GC better than FC), and computational efficiency (FC better than GC). The proposed new algorithm combines the best of both GC and FC and achieves an intermediate speed close to that of FC. In Section 2 we introduce the elements of the GGC framework needed in the rest of the paper. Following the new algorithm description in Section 3, experimental results are shown in Section 4, and our conclusions are summarized in Section 5.

The new algorithm, GC_{sum}^{max}, as well as its theoretical properties (see Section 3.3) were previously announced in (Ciesielski et al., 2012a). However, this conference publication contained neither theoretical (proof of Theorem 3.2) nor extensive experimental justification of the advantages of GC_{sum}^{max} we show here.

2. Terminology: the GGC framework

A *digital image* will be identified with a pair $I = \langle C, f \rangle$, where C is its *domain* (whose elements will be referred to as *spels*, short for space elements) and $f : C \to \mathbb{R}^\ell$ is the image *intensity* (or *attribute*) *function*, with the value $f(c)$ of f at c representing image intensity (an ℓ-dimensional vector, each component of which indicates a measure of some aspect of the signal, like tissue property or color) at the spel c.

It is assumed that the image domain comes with an *adjacency relation*, which decides which pairs of spels are adjacent. An image domain C together with its adjacency structure is referred to as a *scene*.

In a given input image $I = \langle C, f \rangle$, an object $P \subset C$ is identified with its characteristic function χ_P, that is, a mapping $x: C \to \{0, 1\}$ such that $x(c) = 1$ if, and only if, $c \in P$. We usually restrict the collection $\mathcal{X}$ of all allowable objects by indicating two disjoint sets of spels, referred to as *seeds*: $S \subset C$ indicating spels in the *object* and $T \subset C$ pointing out to spels from the *background*. This restricts the collec-

tion of plausible outputs of the algorithm to $\mathcal{X}(S, T) = \{\chi_P : S \subset P \subset C \setminus T\}$.

In every algorithm within GGC, a digital image $I = \langle C, f \rangle$ is identified with a weighted directed graph $G = \langle V, E, w \rangle$ having the following properties. V is the set of vertices of the graph and is equal to the image domain C. (In the case of GC, the set V is often expanded by two additional "virtual" vertices, a source and a sink.) E is set of all edges $\{c, d\}$ in G, that is, the pairs $\langle c, d \rangle$ (identified with $\langle d, c \rangle$) such that the spels c and d are adjacent. $w: E \to [0, \infty)$ is a weight function associating with any edge $e \in E$ its weight $w(e)$.

For $q \in [1, \infty)$ consider the energy functional $\varepsilon_q : \mathcal{X} \to [0, \infty)$, where, for every $x \in \mathcal{X}$, $\varepsilon_q(x)$ is defined as the q-norm of the functional $F_x : E \to \mathbb{R}$ given by a formula $F_x(c, d) = w(c, d)|x(c) - x(d)|$ for $\langle c, d \rangle \in E$. In particular for $q = 1$ and $q = \infty$ we have, $\varepsilon_1(x) = \|F_x\|_1 = \sum_{\langle c,d \rangle \in E} w(c, d)|x(c) - x(d)|$ and $\varepsilon_\infty(x) = \|F_x\|_\infty = \max_{\langle c, d \rangle \in E} w(c, d)|x(c) - x(d)|$.

For $1 \leqslant q \leqslant \infty$, a fixed weighted graph G, and seed sets S and T, let $\varepsilon_{\min}^q$ be the minimum of the energy $\varepsilon_q(x)$ over all objects $x \in \mathcal{X}(S, T)$; that is, $\varepsilon_{\min}^q = \min\{\varepsilon_q(x) : x \in \mathcal{X}(S, T)\}$. Any element of the family $\mathcal{X}_q(S, T) = \{x \in \mathcal{X}(S, T) : \varepsilon_q(x) = \varepsilon_{\min}^q\}$ will be referred to as an energy ε_q minimizer of $\mathcal{X}(S, T)$. Any algorithm A that, given an image I and seed sets S and T, returns an object $A(I, S, T)$ from $\mathcal{X}_q(S, T)$ will be referred to as an ε_q-*minimizing algorithm*.

The standard min-cut/max-flow algorithm is an ε_1-minimizing algorithm. We will use a symbol GC_{sum} to denote this algorithm. We have recently proved (Ciesielski and Udupa, 2011b) (see also (Ciesielski et al., 2011, 2012b)) that both *Relative Fuzzy Connectedness, RFC*, and *Iterative Relative Fuzzy Connectedness, IRFC*, algorithms are the ε_∞-minimizing algorithms. Moreover, we proposed in (Ciesielski et al., 2012b) an IRFC segmentation algorithm, which we named GC_{max}, based on the Optimum Path Forest Framework (Falcão et al., 2004), and proved that it runs in linear time with respect to the image size.

3. The new algorithm

3.1. Motivation

The ε_1- and ε_∞-minimizing algorithms, GC_{sum} and GC_{max}, have their complementary strengths and weaknesses (Ciesielski and Udupa, 2011b; Ciesielski et al., 2011, 2012b). From the point of view of this paper, the most important differences between these algorithms lie in the sensitivity of their output to the choice of the seed sets and the nature of the object's boundary in the input image. Specifically,

Choice of seeds. The outcome of the ε_∞-minimizing algorithm, GC_{max} as well as its older versions RFC and IRFC, are completely unaffected by any changes of the seed sets within the delineated object. (For specifics, see Ciesielski and Udupa (2011b, thm 4) or Theorem 3.1 below.) In particular, relatively small sets of seeds, chosen with little care, lead to the same output as carefully chosen seeds, close to the actual final delineated object.

On the other hand, the ε_1-minimizing algorithms, GC_{sum} and its more effective versions, are highly sensitive to the choice of seeds. This behavior, known as the shrinking problem, is especially acute when the sets of seeds are small, in which case the algorithm has a tendency to output, as a delineated object, a small set very close to the set of object-indicating seeds.[1] Therefore, the GC_{sum} algorithm

---

[1] The shrinking problem has been addressed by many authors, via modifications of the GC method. The best known among these modifications is the method of normalized cuts (see (Shi and Malik, 2000)), in which the energy ε^{sum} is replaced by another "normalized" measure of energy cost, similar to, but different from the ε_1 energy. Moreover, finding the resulting delineation minimizing this new energy measure is NP-hard (see (Shi and Malik, 2000)), and so only approximate solutions can be found in practical time.

requires a careful choice of seeds, relatively close to the desired object boundary, in order for the actual delineated object to be close to the desired object. Such a careful seed choice is especially difficult to achieve automatically in 3D images and to reproduce, if the choice is made by an operator.

Poorly defined boundary. The ε_1-minimizing algorithms, including the GC_{sum} algorithm, have a tendency to choose the objects with small boundary. Although this may lead to an object shrinking problem, this is not an issue, when the input seed sets are relatively large, especially, when they are relatively close to the desired boundary of object and background. At the same time, the tendency of choosing the objects with small boundaries decreases the chance that an output object crosses a true weakly visible boundary or a gap in a boundary, so it reduces the likelihood of causing delineation errors, usually referred to as leaking problems. Moreover, this decreasing of boundary size has a boundary smoothing effect, a feature that may be desirable in certain image segmentation tasks.

On the other hand, the output of the ε_∞-minimizing algorithms is independent of the object boundary size. So, their output has a greater chance of being scraggly and/or passing through gaps in the true object boundary.

3.2. The algorithm

To combine the strengths of both kinds of minimization strategies, we devised the following algorithm. Basically, we obtain a first approximation of the object by applying the GC_{max} algorithm with a conservative weight function, so leakage is minimized. We obtain the final delineation by applying GC_{sum} to the output of thus created first approximation. The first step, optimizing ε_∞, enlarges the initially specified (possibly small) sets of seeds, preserving the algorithm's robustness (with respect to seed choice) and avoiding the shrinking problem of GC_{sum}. The second step, optimizing ε_1, refines this approximation by enlarging it further to an object with a smoother boundary. This final increase creates only a small risk of object leakage, the attribute of GC_{sum}.

More specifically, the algorithm, called[2] GC_{sum}^{max}, is as follows.

Algorithm GC_{sum}^{max}
Input: An image $I = \langle C, f \rangle$ and non-empty disjoint sets:
 $S \subset C$ indicating object, $T \subset C$ indicating background.
Output: An object χ_P from $\mathcal{X}(S, T)$.
begin
  1. create the weighted graph $G = \langle V, E, w \rangle$ associated with I;
  2. use the RFC version of GC_{max} on G to find the sets:
     $\widehat{S}$ with $\chi_{\widehat{S}} \in \mathcal{X}_\infty(S, T)$ and $\widehat{T}$ with $\chi_{\widehat{T}} \in \mathcal{X}_\infty(T, S)$;
  3. apply GC_{sum} to G using $\widehat{S}$ and $\widehat{T}$ as new seed sets to find χ_P;
  4. return χ_P;
end

The fact that both algorithms, GC_{max} and GC_{sum}, can use the same weighted graph G, associated with the input image I, makes the merging of these two algorithms seamless and effortless.

Line 1 of the algorithm constitutes an "implicit parameter" of the algorithm, as mentioned above. The choice of the weight function, which in FC literature is called the *affinity* function, is explained in more detail in the next section.

Line 3 is straightforward. Since the sets $\widehat{S}$ and $\widehat{T}$ of spels, output by the GC_{max} step, are typically already quite large and close to the

desired object boundary, there is little danger of shrinkage. Also, GC_{sum} has a smoothing effect on the final output.

Line 2 requires a few words of explanation. To find $\widehat{S}$, we run GC_{max} in a version described in Ciesielski et al. (2012b, Section 4.3) which, in particular, returns a function $\mu^C(c, W)$, of variable c from $C = V$ into $[0,1]$, denoting the strength of connectedness of spel $c \in C$ to the set W of spels. We run GC_{max} twice, once with $W = S$ and once with $W = T$, calculating functions $\mu^C(c, S)$ and $\mu^C(c, T)$, respectively. The RFC object $\widehat{S}$ is simply defined as the set $\{c \in C: \mu^C(c, S) > \mu^C(c, T)\}$. Similarly, the RFC co-object is defined as $\widehat{T} = \{c \in C : \mu^C(c, T) > \mu^C(c, S)\}$. Since GC_{max} runs in a linear time with respect to the image size $|C|$, a fact theoretically proved in (Ciesielski et al., 2012b), this does not add much to a total run time of the algorithm, especially in comparison with the running time of the GC_{sum} component, which runs in time of order $O(|C|^{2.5})$ or greater. We choose the RFC segmented objects $\widehat{S}$ and $\widehat{T}$ as the new seed sets rather than their IRFC counterparts—the standard output of GC_{max}—since they are smaller (see Ciesielski et al., 2012b, Theorem 4.3(iii)), while they still belong to $\mathcal{X}_\infty(S, T)$ and $\mathcal{X}_\infty(T, S)$, respectively, insuring $\mathcal{X}_\infty(\widehat{S}, \widehat{T}) = \mathcal{X}_\infty(S, T)$. This leaves some extra room for the GC_{sum} step of the algorithm to act upon, which chooses an object from $\mathcal{X}(\widehat{S}, \widehat{T})$, while preserving the extrema choices, $\widehat{S}$ and $\widehat{T}$, indicated by GC_{max}. (Note that $\mathcal{X}_1(\widehat{S}, \widehat{T})$ need not be equal to $\mathcal{X}_1(S, T)$. In fact, $\mathcal{X}_1(\widehat{S}, \widehat{T})$ can be disjoint with $\mathcal{X}_\infty(\widehat{S}, \widehat{T}) = \mathcal{X}_\infty(S, T)$.)

It is interesting to observe that in GC_{sum}^{max}, the GC_{sum} component always runs in a reduced set $D = \{c \in C: \mu^C(c, T) = \mu^C(c, S)\}$, where RFC finds a tie in strength of connectedness with respect to seed sets S and T. This makes the final execution time of GC_{sum}^{max} to be $O(|C| + |D|^{2.5})$, where $|D|$ is typically proportional to the size of the boundary. In practical situations, wherein object boundary size is much (10–80 times) smaller than its volume, $|D| \ll |C|$, and so the execution time of GC_{sum}^{max} is quasi-linear. Note also that sequences of executing GC_{sum} and GC_{max}, such as $GC_{sum} \to GC_{max}$, $GC_{sum} \to GC_{max} \to GC_{sum}$, and $GC_{max} \to GC_{sum} \to GC_{max}$, are not useful from the viewpoint of the goal of this paper.

Fig. 1 shows, via an example, what happens in Steps 2 and 3. The RFC objects, shown in (d) and (e), represent, respectively, the sets $\widehat{S}$ and $\widehat{T}$. The output of GC_{sum}, by employing $\widehat{S}$ and $\widehat{T}$ as seed sets, is shown in (f). This constitutes the output of the GC_{sum}^{max} algorithm, which should be compared to the outputs of: (b) IRFC, (c) Graph Cut GC_{sum}, and (d) RFC algorithms.

### 3.3. Nice properties of GC_{sum}^{max}

Robustness with respect to seed set size and location. We have the following theorem, which immediately follows from a similar result on RFC segmentations. (See e.g. (Ciesielski and Udupa, 2011b; Ciesielski et al., 2007).)

Theorem 3.1. *Let $I = \langle C, f \rangle$ be an image and $S, T \subset C$ non-empty disjoint sets of seeds. If $S' \subset \widehat{S}$ intersects every connected component of $\widehat{S}$ and $T' \subset \widehat{T}$ intersects every connected component of $\widehat{T}$, then $GC_{sum}^{max}(I, S, T)$ and $GC_{sum}^{max}(I, S', T')$ have identical outputs.*

In particular, if each of $\widehat{S}$ and $\widehat{T}$ has only one connected component in the graph G, then any other choice of non-empty sets of seeds $S' \subset \widehat{S}$ and $T' \subset \widehat{T}$ leads to identical delineations. Note in particular that, in this case, even if S' and T' are singleton sets, the identity of outputs is guaranteed.

Robustness with respect to remapping the image intensity by an increasing function. The validity of this property depends very much on the way the weight/affinity function is created from the image intensity function. We have the following results concerning the *weight* function.

---

[2] The chosen name stresses that we use, consecutively, the optimizers of max and sum energies. Another appealing name would be RFC-GC, to stress more conventional names of the algorithmic components.

Fig. 1. (a) A region of interest around a tendon in an MRI slice of the wrist. The tendon cross section appears dark. S consists of one spel indicated inside the object by a +mark. T consists of two spels indicated by the other two +marks in the surrounding tissue. (b) IRFC object (obtained by the full version of GC_{max}) using 8-neighborhood with the cost function $w(c, d) = K − |f(c) − f(d)|$. (c) GC_{sum} object collapses to just the seed set S. (d) RFC object $\widehat{S}$ for the internal seed set S. (e) RFC object $\widehat{T}$ for the external seed set T. (f) The GC_{sum}^{max} returned object, resulting from applying GC_{sum} to sets $\widehat{S}$ (from (d)) and $\widehat{T}$ (from (e)) used as seeds.

Theorem 3.2. *Let $I = \langle C, f \rangle$ and $I' = \langle C, f' \rangle$ be the images with associated weighted graphs $G = \langle V, E, w \rangle$ and $G' = \langle V, E, w' \rangle$, respectively. If w' is a modification of w via an increasing linear function (i.e., if w' is a composition $L \circ w$ of w and a linear function L), then for every seed sets $S, T \subset C$, the outputs of $GC_{sum}^{max}(I, S, T)$ and $GC_{sum}^{max}(I', S, T)$ are identical.*

More generally, if w' is a modification of w via an increasing function, then the associated RFC approximations $\langle \widehat{S}, \widehat{T} \rangle$ and $\langle \widehat{S}', \widehat{T}' \rangle$ are identical.

Proof. Let the weight function w' be a modification of w via an increasing function h, that is, $w' = h \circ w$. Then the resulting RFC approximations $\langle \widehat{S}, \widehat{T} \rangle$ and $\langle \widehat{S}', \widehat{T}' \rangle$ are identical, as proved in Ciesielski and Udupa (2010, Section 2). If, moreover, function h is linear (i.e., $h(z) = az + b$ so that $w'(c, d) = h(w(c, d)) = a\, w(c, d) + b$), then the final GC_{sum} outputs are also identical, since, in this situation, for every $x_1, x_2 \in \mathcal{X}$,

$$\sum_{\langle c,d \rangle \in E} w(c, d)|x_1(c) - x_1(d)| \leqslant \sum_{\langle c,d \rangle \in E} w(c, d)|x_2(c) - x_2(d)|$$

if, and only if,

$$\sum_{\langle c,d \rangle \in E} w'(c, d)|x_1(c) - x_1(d)| \leqslant \sum_{\langle c,d \rangle \in E} w'(c, d)|x_2(c) - x_2(d)|.$$

So, the inequality $\varepsilon_1(x_1) \leqslant \varepsilon_1(x_2)$ is true for the weight function w if, and only if, it is true for w'.   □

Notice that for many methods of constructing weight/affinity functions w from the intensity map f (Saha et al., 2000; Ciesielski and Udupa, 2010) (e.g., for the homogeneity based affinity), the linear modification $h(z) = az + b$ of the intensity function with $0 < a < 1$ translates into an increasing linear modification of w.

Other nice properties. The following properties are difficult to express in the formal mathematical language, as the theorems above. Nevertheless, they can be argued at a less formal level and demonstrated empirically. Some level of *boundary smoothness* is assured for the output of GC_{sum}^{max} by a similar property of GC_{sum}. Similarly, some level of *leakage control* is achieved. The greater robustness

(insensitivity) to the artifacts such as a *slow background variation* component modulating the image intensity function can be achieved by a careful creation of w, a property stemming from FC as demonstrated in (Udupa and Samarasekera, 1996).

4. Experimental results

In this section we present results for accuracy and efficiency of segmentation based on experiments carried out on three sets of data: 2D medical images (without and with added noise), 2D natural color images, and 3D medical images. True segmentations for the medical images were created by manual or user-steered but algorithm-determined (Falcão et al., 1998) outlining by experts knowledgeable in the domain. These data have been used in several earlier publications, e.g. in (Miranda et al., 2009) and in (Rother et al., 2004). For the medical images true segmentations were created by experts on a slice-by-slice basis. For the natural images, we used the true segmentations available on line.

4.1. 2D medical images

These data comprise of 40 MR image slices of the human foot. The 40 slices were selected from 3D MRI scene data pertaining to the left or the right foot of 20 different live subjects. Several (1–3) slices were selected randomly approximately at similar anatomic locations from each 3D scene. The objects of interest in these slices are the two large bones in the peritalar complex, namely, the talus and the calcaneus. These data were chosen since they present real and practical situations for some of the issues mentioned earlier (leakage, shrinkage, etc.). Cortical bones elicit very little signal and so do connective tissues such as ligaments and tendons, yet the objects we seek have smooth boundaries. Fig. 2(a) demonstrates the context of the two objects in a sample image.

We compared four different algorithms in the task of segmenting calcaneus and talus: the linear time IRFC algorithm, implemented as GC_{max} (Ciesielski et al., 2012b); the Graph Cut GC max-flow algorithm GC_{sum}; our new algorithm GC_{sum}^{max}; and the power watershed PW algorithm $PW_{q=2}$ (Couprie et al., 2011). The

Fig. 2. (a) An example of an MR image slice of a foot from our 2D medical image data set showing talus and calcaneus bones. (b) True segmentation of the talus used to create two examples of seed sets, the curves shown in (c) and (d), for the object and background obtained by eroding the true segmentation.

GC$_{sum}$ code comes from a software library in C++ developed by Yuri Boykov and Vladimir Kolmogorov. It implements the max-flow algorithm as described in (Boykov and Kolmogorov, 2004). The GC$_{sum}^{max}$ implementation combines the codes of the RFC version of GC$_{max}$ algorithm, as described in (Ciesielski et al., 2012b), with GC$_{sum}$. The power watershed algorithm (Couprie et al., 2011) is a recently introduced version of the IRFC algorithm GC$_{max}$, as discussed in detail in (Miranda and Falcão, 2009). The PW$_{q=2}$ code comes from a software library in C developed by Camille Couprie, which is available at sourceforge: http://sourceforge.net/projects/powerwatershed/. The experiments presented in this section were conducted on an AMD Athlon 64 X2 Dual-Core Processor TK-57 (1.9 GHz, 2 × 256 KB L2 cache) with 2 GB of RAM.

The arc weights $w(c, d)$ were computed as the complement of the difference of image intensities (i.e., as $K - |f(c) - f(d)|$, where K stands for the maximum value of $|f(c) - f(d)|$ in the scene), see (Miranda et al., 2010) for more complex weight functions. To assess segmentation accuracy, we used Dice coefficient.

4.1.1. Seeds chosen by erosion

In these experiments, the seeds were chosen by erosion of different magnitude, that is, the seeds constitute the boundaries of the foreground and background objects eroded in a specified way, see e.g. (Sinop and Grady, 2007). This allows varying the seed set in a controlled manner compared to the alternative of operators specifying seeds interactively, and thereby we can study the influence of seed sets on results also in a controlled manner.

Mean accuracy estimated over the 40 images is plotted in Figs. 3(a) and 4(a) for talus and calcaneus as a function of varying seed set size for the four methods. Figs. 3(b) and 4(b) similarly display the average curves of computational time for the four

methods. The rank ordering of these methods over all images for the accuracy is depicted in Fig. 5.

From Figs. 3–5 we observe that GC$_{sum}^{max}$ outperformed the remaining three algorithms. Similarly, the accuracy of GC$_{sum}$ was consistently lower than that of the other three algorithms. On the other hand, the distinction between PW$_{q=2}$ and GC$_{max}$ is not apparent. While for calcaneus, Fig. 4(a), PW$_{q=2}$ presented better results than GC$_{max}$, for talus the results of PW$_{q=2}$ and GC$_{max}$, Fig. 3(a), are quite similar.

The lack of a clear cut distinction between PW$_{q=2}$ and GC$_{max}$ is in agreement with our theoretical result of Ciesielski et al. (2012b, Section 4) that, under the assumption that there are no tie zones (i.e., spels having the same strength of connectedness values $\mu(\cdot, S)$ and $\mu(\cdot, T)$ with respect to the object and background seeds), the outputs of PW$_{q=2}$ and GC$_{max}$ are identical. In this light, the similarity of the talus segmentations returned by PW$_{q=2}$ and GC$_{max}$ can be interpreted as a lack of large tie zones for this object. At the same time, the differences between calcaneus segmentations of PW$_{q=2}$ and GC$_{max}$ were caused, most likely, by the existence of larger tie zones for the calcaneus.

Finally, notice that, according to our theoretical results from Ciesielski et al. (2012b, Section 4), the replacement of the graph weight function w with w^q (i.e., substituting $A_{l \to w^q}$ for $A_{l \to w}$) for the large values of q makes the outputs of GC$_{max}$ and GC$_{sum}$ (and so, also of PW and GC$_{sum}^{max}$) essentially identical.

In regard to computational time, Figs. 3(b) and 4(b), it was not possible to measure the exact running time of PW$_{q=2}$, because its source code does not allow us to restrict the computation to only inside the band generated by erosion (i.e., the band between the internal and external seeds). It seemingly executed over the entire image graph taking more time than necessary, so it took an almost
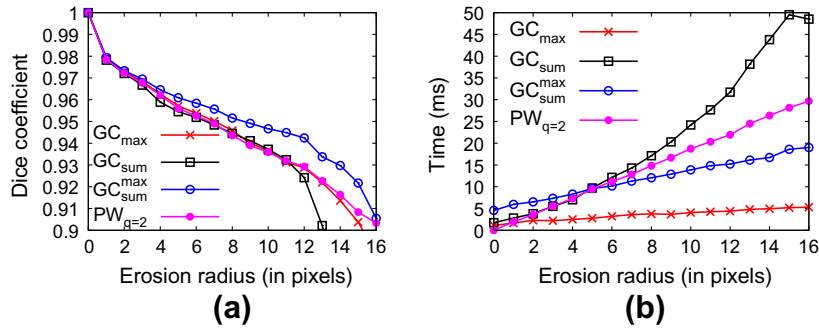
Fig. 3. (a) The mean accuracy curves and (b) running times for different methods for the segmentation of talus.
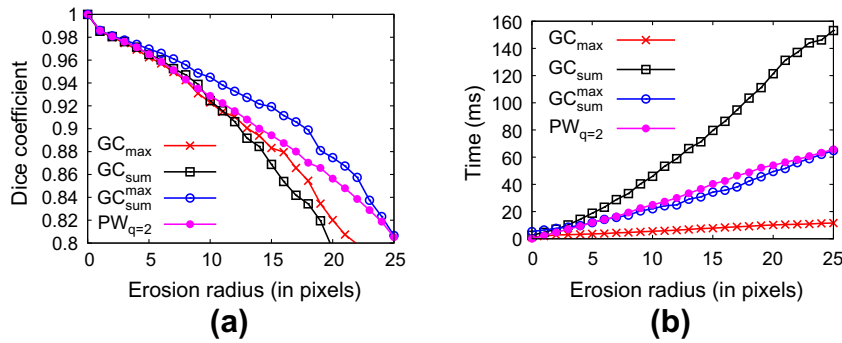

Fig. 4. (a) The mean accuracy curves and (b) running times for different methods for the segmentation of calcaneus.


Fig. 5. For each individual image, the methods may be ranked according to their mean Dice coefficient values, as first (best), second, third, or fourth (worst). By computing the frequency for each rank position, we have a rank distribution for segmenting: (a) talus and (b) calcaneus.

constant and large running time for any erosion level. However, it was possible to measure an estimated running time by normalization of its measured running time, taking into account the relative size of the band in relation to the image domain size. In fact, this gives us an upper-bound for its actual running time inside the band, and this is what is depicted in Figs. 3(b) and 4(b).

In Fig. 6 we display images of some sample segmentations for the talus.

4.1.2. 2D medical images with added noise

To evaluate the robustness of the methods in relation to different noise levels, we conducted experiments on the same set of 40 2D images with different amounts of added white noise. The random noise was obtained using the ImageMagick command-line tools. Fig. 7 shows some sample noise-corrupted images used in our experiments. The accuracy curves for the segmentation of talus and calcaneus are given in Figs. 8 and 9. Again a relative behavior

similar to that observed in Figs. 3 and 4 is demonstrated by the algorithms under added noise.

4.1.3. Seeds chosen by a robot user

In these experiments, the seeds were chosen by a robot user, the method introduced by (Gulshan et al., 2010), to simulate user interaction of interactive segmentation by placing brush strokes automatically to iteratively, and interactively, complete the segmentation task. The procedure, including the choice of seeds (based on the knowledge of the ground truth) and the successive delineations, is iterative. The initial seeds for the object are placed at the point(s) in the object farthest from the boundary; similarly, for the background. At each successive iteration, the robot always places a circular brush stroke in the largest connected component of the segmentation error area (placed at the point(s) farthest from the boundary of the component), and updates the segmentation. The process is repeated up to 20 times, generating a sequence of 20 simulated user strokes.

Fig. 6. Some sample results for the talus for the seed sets depicted in Figs. 2(b) (top row) and in 2(c) (bottom row): (a and d) GC_{sum}, (b and e) GC_{max}, (c and f) GC_{sum}^{max}. Compare these with the true segmentation in Fig. 2(a).

Fig. 7. Examples of images with random (additive and white) noise. (a) Original image, (b) original with 20% noise, and (c) original with 30% noise.
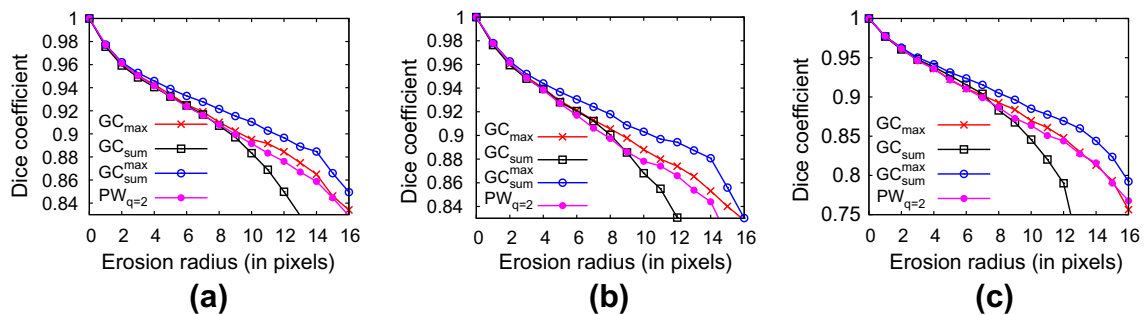
Fig. 8. The mean accuracy curves for different methods for the segmentation of talus in the images with different noise levels: (a) 10%, (b) 20%, and (c) 30%.

We used a circular brush with a diameter of 17 pixels for the grabcut dataset, as suggested in (Gulshan et al., 2010). For the talus and calcaneus, we used a smaller diameter of 10 pixels, in order to be more compatible with the smaller image size of this dataset (256×256 pixels). Fig. 10 shows the experimental curves. We observe again that GC_{sum}^{max} outperforms other algorithms. Interestingly, GC_{sum} tends to approach the performance of GC_{sum}^{max}, as do other algorithms. This demonstration reemphasizes the need for

carefully specified seed sets for GC_{sum} in order to reap good performance from it.

4.2. 2D natural color images

A dataset with 50 natural color images with known true segmentations was obtained from (Martin et al.; Rother et al.) for these experiments. The pixels in the reference segmentations
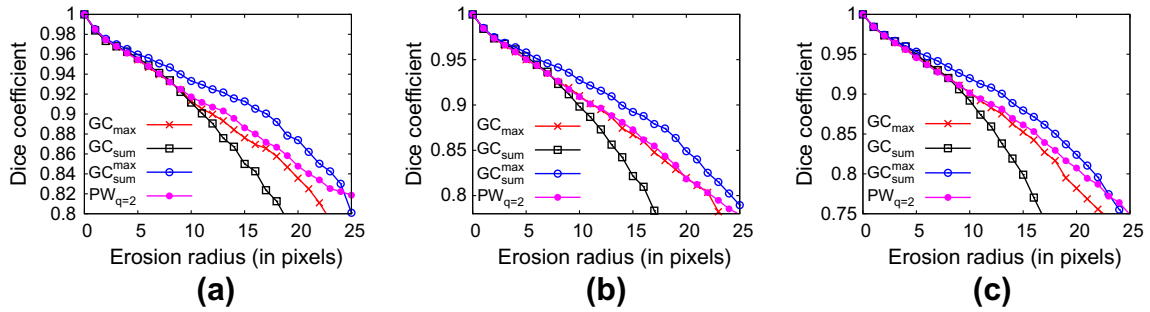
Fig. 9. The mean accuracy curves for different methods for the segmentation of calcaneus in the images with different noise levels: (a) 10%, (b) 20%, and (c) 30%.
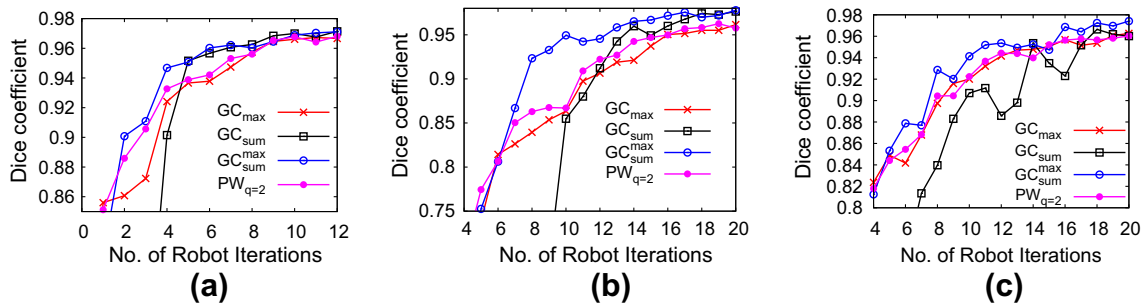

Fig. 10. Results using a robot user for segmenting: (a) talus, (b) calcaneus, and (c) grabcut dataset.


Fig. 11. (a) The mean accuracy curve and (b) running times for different methods considering the mixed areas as object for the color image data set.

represent three different regions: background, object, and mixed areas. The last category represents a set of pixels which cannot be unequivocally allocated to either background or object. Since we are interested only in one object, we conducted two separate experiments, the first considering the mixed areas as object and the second taking the mixed areas as background. The goal of these experiments was to show that the final conclusions were not affected by this choice. The arc weights $w(c, d)$ were computed as the complement of the maximum difference of intensity values over the three channels, that is, given by a formula $K - \max\{|f_R(c) - f_R(d)|, \quad |f_G(c) - f_G(d)|, \quad |f_B(c) - f_B(d)|\}$, where $f(c) = \langle f_R(c), f_G(c), f_B(c) \rangle$. For more on how to chose suitable weights see (Miranda et al., 2010). The experiments presented in this subsection were conducted on an Intel Core i7-2630QM CPU at 2.00 GHz × 8, with 8 GB of RAM.

The resulting mean curves are displayed in Figs. 11 and 12, and some sample images with results are displayed in Fig. 13 for different methods. The graphs show that the combined approach GC_{sum}^{max} provided slightly better results. In the case of heterogeneous databases, such as the 2D color images used in this experiment, there

is no best method in the absolute sense for all the images, because each image has its own characteristics which may favor a particular approach.

To study this phenomenon, we present also the rank distributions in Fig. 14, obtained by comparing the methods for each individual image separately. It is readily seen that the proposed method was the best for most of the images. Fig. 15 displays sample results, similar to that shown in Fig. 13, but for a user selected set of seeds.

4.3. 3D medical images

In the first experiment, we used a dataset of 3D MRI brain images of 20 normal subjects. We performed the 3D segmentation of the cerebellum for all methods in all data sets. Segmentation and separation of the cerebellum from the rest of the brain region is crucial in many neurological applications, such as the study of schizophrenia (Puri et al., 1996) and epilepsy (Hagemann et al., 2002). T1-weighted images were acquired on a 2T Elscint scanner and at a voxel size of $0.98 \times 0.98 \times 1.00 \text{ mm}^3$. The arc weights

Fig. 12. (a) The mean accuracy curve and (b) running times for different methods considering the mixed areas as background for the color image data set.


Fig. 13. Results with seeds obtained by eroding and dilating the ground-truth. First column: Input images and seeds shown with solid colors. Second column: GC$_{max}$ results (more specifically IRFC, but PW outputs similar results). Third column: GC$_{sum}$ results. Fourth column: GC$_{sum}^{max}$ results for the same seeds.

$w(c, d)$ were computed as the complement of the difference of image intensities (i.e., $K - |f(c) - f(d)|$) using 6-neighbors for spel adjacency. The experiments were carried out on an Intel Core i7-2630QM CPU at 2.00 GHz × 8, with 8 GB of RAM.

Figs. 16 and 17 show the obtained results. Again, as demonstrated in various 2D images, the comparative behavior of the four algorithms is borne out in this 3D segmentation task as well.

In the last experiment, we used a 3D dataset, composed of 20 MR images of the foot. The objects of interest in these 3D volumes are the two large bones in the peritalar complex, namely, the talus and the calcaneus. The image parameters are the same as for the 2D foot data; the slice spacing was 1.5 mm. We computed the mean accuracy curve (over the 20 3D scenes) for the four methods to segment the talus and the calcaneus, as a function of the

Fig. 14. For each individual image in the color image data set, the methods may be ranked according to their mean Dice coefficient values, as first (best), second, third, or fourth (worst). By computing the frequency for each rank position, we have a rank distribution: (a) considering the mixed areas as object and (b) considering the mixed areas as background.

Fig. 15. Examples with user-selected seeds. First column: Input images and seeds shown with solid colors. Second column: GC_{max} results (more specifically IRFC, but PW outputs similar results). Third column: GC_{sum} results. Forth column: GC_{sum}^{max} results.
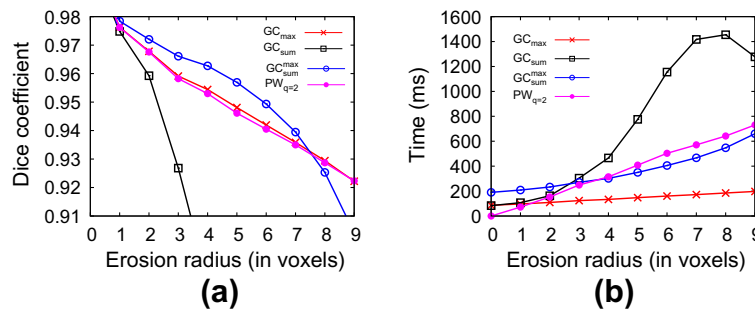

Fig. 16. (a) Mean accuracy curves and (b) running times for different methods for segmenting the cerebellum.

different seed sets obtained by eroding the true segmented objects. The results are shown in Fig. 18.

5. Conclusions

The focus of research in this paper was on pI approaches, particularly graph-based combinatorial optimization techniques. As we noted in Section 1, there are two contradictory requirements to be fulfilled by any pI-based segmentation algorithm—gathering the object parts no mater how compact or how scattered the object is. Some mathematical constraints, such as those based on GC max-flow, inherently favor compactness. Others, such as FC, have an inherent bias to favor scattered objects. The shrinking problem apart, the former strategies often fail to gather the loose and detailed peripheral parts of the same object. The latter strategies, conversely, often get misled and collect non-object surrounding regions considering them as loose aspects of the same object. It is a challenge to balance off these opposing requirements within a

Fig. 17. Sample 3D results for segmenting the cerebellum by using different algorithms: (a–d) GC_{sum}, (e–h) GC_{max}, and (i–l) GC_{sum}^{max}. Seeds were generated by using an erosion radius of 4 voxels.

Fig. 18. (a) The mean accuracy curves and (b) running times for different methods for the 3D segmentation of talus. (c) The mean accuracy curves and (d) running times for different methods for the 3D segmentation of calcaneus.

single mathematical framework. (We note that often the same object has both compact and loose parts, which precludes the possibility of choosing a strategy that best suits the object characteristics.) This paper makes an important contribution, we believe, by combining within a single framework two methods, GC and RFC, representative of the above two strategies, for achieving an optimal trade-off. It also demonstrates how some desirable properties (such as robustness to seed set size and location, as well as speed) are maintained by the new strategy. Our empirical evaluation on a variety of images indicates that the new strategy consistently outperforms other state-of-the-art pI techniques such as PW and IRFC.

There are some avenues for possibly further advancing the proposed approach. The new strategy seems to be best suited within a hybrid (pI + SM) approach wherein a prior object model (in applications where it is feasible to create a model) can guide conservatively the automatic selection of a small set of seeds. Another possible extension is generalizing the framework from two objects

(object and background) to many objects. While RFC itself poses no problem to this effort (Ciesielski et al., 2007), the GC part of the iteration is more difficult to handle. The additional constraints provided by the large seed sets (segmentations, which get produced by RFC) may make this problem more tractable than the multi-object GC_{sum} segmentation problem.

References

Angelini-Casadevall, E.D., Imielinska, C., Jin, Y., Laine, A.F., 2002. Improving statistics for hybrid segmentation of high-resolution multichannel images. In: Proc. SPIE Int. Soc. Opt. Eng. 4684.

Aubert, G., Blanc-Féraud, L., 1999. Some remarks on the equivalence between 2d and 3d classical snakes and geodesic active contours. International Journal of Computer Vision 34, 19–28.

Bai, X., Sapiro, G., 2007. Distance cut: interactive segmentation and matting of images and videos. In: Proceedings of the IEEE International Conference on Image Processing (ICIP), San Antonio, Texas, USA, pp. 249–252.

Boykov, Y., Kolmogorov, V., 2004. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. IEEE Transactions on Pattern Analysis and Machine Intelligence 26, 1124–1137.

Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. IEEE Transactons on Pattern Analysis and Machine Intelligence 23, 1222–1239.

Ciesielski, K.C., Miranda, P.A.V., Udupa, J.K., Falcão, A.X., 2012a. Image segmentation by combining the strengths of relative fuzzy connectedness and graph cut. In: Proceedings of the IEEE International Conference on Image Processing (ICIP), pp. 2005–2008.

Ciesielski, K.C., Udupa, J.K., 2010. Affinity functions in fuzzy connectedness based image segmentation I: equivalence of affinities. Computer Vision and Image Understanding 114, 146–154.

Ciesielski, K.C., Udupa, J.K., 2011a. A framework for comparing different image segmentation methods and its use in studying equivalences between level set and fuzzy connectedness frameworks. Computer Vision and Image Understanding 115, 721–734.

Ciesielski, K.C., Udupa, J.K., 2011b. Region-based segmentation: fuzzy connectedness, graph cut, and other related algorithms. In: Deserno, T.M. (Ed.), Biomedical Image Processing. Springer-Verlag.

Ciesielski, K.C., Udupa, J.K., Falcão, A.X., Miranda, P.A.V., 2011. Comparison of fuzzy connectedness and graph cut segmentation algorithms. In: Medical Imaging: Image Processing, SPIE Proceedings.

Ciesielski, K.C., Udupa, J.K., Falcão, A.X., Miranda, P.A.V., 2012b. Fuzzy connectedness image segmentation in graph cut formulation: a linear-time algorithm and a comparative analysis. Journal of Mathematical Imaging and Vision 44, 375–398.

Ciesielski, K.C., Udupa, J.K., Saha, P.K., Zhuge, Y., 2007. Iterative relative fuzzy connectedness for multiple objects, allowing multiple seeds. Computer Vision and Image Understanding 107, 160–182.

Cootes, T., Edwards, G., Taylor, C., 1999. Comparing active shape models with active appearance models. In: Pridmore, T., Elliman, D. (Eds.), British Machine Vision Conference, pp. 173–182.

Couprie, C., Grady, L., Najman, L., Talbot, H., 2011. Power watersheds: a unifying graph-based optimization framework. IEEE Transactions on Pattern Analysis and Machine Intelligence 33, 1384–1399.

Falcão, A.X., Bergo, F.P.G., 2004. Interactive volume segmentation with differential image foresting transforms. IEEE Transactions on Medical Imaging 23, 1100–1108.

Falcão, A.X., Costa, L.F., da Cunha, B.S., 2002. Multiscale skeletons by image foresting transform and its applications to neuromorphometry. Pattern Recognition 35, 1571–1582.

Falcão, A.X., Stolfi, J., Lotufo, R.A., 2004. The image foresting transform: theory, algorithms, and applications. IEEE Transactions on Pattern Analysis and Machine Intelligence 26, 19–29.

Falcão, A.X., Udupa, J.K., Samarasekera, S., Sharma, S., Hirsch, B.E., Lotufo, R., 1998. User-steered image segmentation paradigms: live wire and live lane. Graphical Models and Image Processing 60, 233–260.

Grady, L., 2006. Random walks for image segmentation. IEEE Transactions on Medical Imaging 28, 1768–1783.

Gulshan, V., Rother, C., Criminisi, A., Blake, A., Zisserman, A., 2010. Geodesic star convexity for interactive image segmentation. In: Proc. of Computer Vision and Pattern Recognition, pp. 3129–3136.

Hagemann, G., Lemieux, L., Free, S.L., Krakow, K., Everitt, A.D., Kendall, B.E., Stevens, J.M., Shorvon, S.D., 2002. Cerebellar volumes in newly diagnosed and chronic epilepsy. Journal of Neurology 249, 1651–1658.

Imielinska, C., Metaxas, D., Udupa, J.K., Jin, Y., Chen, T., 2001. Hybrid segmentation of anatomical data. In: Proceedings of MICCAI, pp. 1048–1057.

Kolmogorov, V., Zabih, R., 2004. What energy functions can be minimized via graph cuts. IEEE Transactions on Pattern Analysis and Machine Intelligence 26, 147–159.

Liang, J., McInerney, T., Terzopoulos, D., 2006. United snakes. Medical Image Analysis 10, 215–233.

Lim, S.J., 2006. Hybrid Method of Fuzzy Connectedness and Active Shape Model Algorithms for Three-dimensional Medical Image Segmentation. Ph.D. Thesis.

Liu, J., Udupa, J.K., 2009. Oriented active shape models. IEEE Transactions on Medical Imaging 28, 571–584.

Malladi, R., Sethian, J., Vemuri, B., 1995. Shape modeling with front propagation: a level set approach. IEEE Transactions on Pattern Analysis and Machine Intelligence 17, 158–175.

Malmberg, F., 2011. Image foresting transform: on-the-fly computation of segmentation boundaries. In: Image Analysis: 17th Scandinavian Conference on Image Analysis. Lecture Notes in Computer Science. Springer, Berlin, pp. 616–624.

Malmberg, F., Lindblad, J., Sladoje, N., Nystrom, I., 2011. A graph-based framework for sub-pixel image segmentation. Theoretical Computer Science 412, 1338–1349.

Malmberg, F., Nystrom, I., Mehnert, A., Engstrom, C., Bengtsson, E., 2010. Relaxed image foresting transforms for interactive volume image segmentation. In: Medical Imaging: Image Processing, SPIE Proceedings.

Martin, D., Fowlkes, C., Tal, D., Malik, J. Berkeley Segmentation Dataset and Benchmark. <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping>.

Miranda, P.A.V., Falcão, A.X., 2009. Links between image segmentation based on optimum-path forest and minimum cut in graph. Journal of Mathematical Imaging and Vision 35, 128–142.

Miranda, P.A.V., Falcão, A.X., Udupa, J.K., 2009. Cloud bank: a multiple clouds model and its use in mr brain image segmentation. In: Proceedings of the Sixth IEEE International Symposium on Biomedical Imaging from Nano to Macro (ISBI), Boston, pp. 506–509.

Miranda, P.A.V., Falcão, A.X., Udupa, J.K., 2010. Synergistic arc-weight estimation for interactive image segmentation using graphs. Computer Vision and Image Understanding 114, 85–99.

Pham, D., 2001. Spatial models for fuzzy clustering. Computer Vision and Image Understanding 84, 285–297.

Pizer, S.M., Gerig, G., Joshi, S., Aylward, S., 2003. Multiscale medial shape-based analysis of image objects. In: Roux, C., Udupa, J.K. (Eds.), Proceedings of the IEEE, pp. 1670–1679.

Protiere, A., Sapiro, G., 2007. Interactive image segmentation via adaptive weighted distances. IEEE Transactions on Image Processing 16, 1046–1057.

Puri, B.K., Davey, N.J., Ellaway, P.H., Lewis, S.W., 1996. An investigation of motor function in schizophrenia using transcranial magnetic stimulation of the motor cortex. The British Journal of Psychiatry 169, 690–695.

Rother, C., Kolmogorov, V., Blake, A., 2004. "grabcut: Interactive foreground extraction using iterated graph cuts. ACM Transactions on Graphics 23, 309–314.

Rother, C., Kolmogorov, V., Blake, A., Brown, M. Image and Video Editing: Grabcut. <http://research.microsoft.com/en-us/um/cambridge/projects/visionimagevideoediting/segmentation/grabcut.htm>.

Saha, P.K., Udupa, J.K., Odhner, D., 2000. Scale-based fuzzy connectedness image segmentation: theory, algorithms, and validation. Computer Vision and Image Understanding 77, 145–174.

Sandor, S., Leahy, R., 1997. Surface-based labeling of cortical anatomy using a deformable atlas. IEEE Transactions on Medical Imaging 16, 41–54.

Shi, J., Malik, J., 2000. Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 22, 888–905.

Singaraju, D., Grady, L., Sinop, A.K., Vidal, R., 2010. Continuous valued mrfs for image segmentation. In: Blake, A., Kohli, P., Rother, C., (Eds.), Advances in Markov Random Fields for Vision and Image Processing.

Sinop, A.K., Grady, L., 2007. A seeded image segmentation framework unifying graph cuts and random walker which yields a new algorithm. In: Proc. of ICCV 07.

Udupa, J.K., Samarasekera, S., 1996. Fuzzy connectedness and object definition: theory, algorithms, and applications in image segmentation. Graphical Models and Image Processing 58, 246–261.