Available online at www.sciencedirect.com

SCIENCE DIRECT®

Discrete Applied Mathematics 151 (2005) 55–77

DISCRETE APPLIED MATHEMATICS

www.elsevier.com/locate/dam

Simultaneous fuzzy segmentation of multiple objects☆

Bruno M. Carvalho[a], Gabor T. Herman[b], T. Yung Kong[c]

[a]*Departamento de Informática e Matemática Aplicada, Universidade Federal do Rio Grande do Norte, Natal, RN 59072-970, Brazil*

[b]*Doctoral Program in Computer Science, Graduate Center - CUNY, New York, NY 10016, USA*

[c]*Department of Computer Science, Queens College, CUNY, Flushing, NY 11367, USA*

Received 9 November 2003; received in revised form 9 May 2004; accepted 10 February 2005
Available online 26 July 2005

Abstract

Fuzzy segmentation is a technique that assigns to each element in an image (which may have been corrupted by noise and/or shading) a grade of membership in an object (which is believed to be contained in the image). In an earlier work, the first two authors extended this concept by presenting and illustrating an algorithm which simultaneously assigns to each element in an image a grade of membership in each one of a number of objects (which are believed to be contained in the image). In this paper, we prove the existence of such a fuzzy segmentation that is uniquely specified by a desirable mathematical property, show further examples of its use in medical imaging, and illustrate that on several biomedical examples a new implementation of the algorithm that produces the segmentation is approximately seven times faster than the previously used implementation. We also compare our method with two recently published related methods.

Keywords: Fuzzy segmentation; Fuzzy graphs; Greedy algorithms

☆ This research has been supported by NIH Grant HL70472 (GTH and BMC) and CAPES-BRAZIL (BMC).

E-mail addresses: motta@dimap.ufrn.br (B.M. Carvalho), gabortherman@yahoo.com (G.T. Herman), ykong@cs.qc.edu (T.Y. Kong).

0166-218X/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.dam.2005.02.031

1. Introduction

Digital image segmentation is the process of assigning distinct labels to different objects in an image. The task of segmenting an object from a background in an image becomes particularly hard for a computer when, instead of the brightness values, what distinguishes the object from the background is some textural property, or when the image is corrupted by noise and/or inhomogeneous illumination. One concept that has been successfully used to achieve segmentation in such corrupted images is fuzzy connectedness, as can be seen in [3,10] and the references therein. Our approach here is a generalization of the one advocated in [15] (based on the work of [12]) to arbitrary digital spaces [7] and simultaneous multiple object segmentation [8]. We present below a proof of the claim in [8] regarding the existence of a segmentation into multiple objects that is uniquely specified by a desirable mathematical property.[1] We also discuss the relationship of the methodology proposed by the first two authors in [8] (and reproduced here) to alternative methods of fuzzy segmentation.

Prior to getting into our theory we give a picturesque description of the approach.

Our model for describing the algorithm takes the form of a military exercise. It involves a number of castles such that there is a one-way road from every castle to every other castle (equivalently, for every pair of distinct castles c and d, there is a one-way road from c to d and a one-way road from d to c). There are also a number of armies. Each road from a castle to another one has an *affinity* for each army; this is measured by a nonnegative integer (the lower this integer, the more difficult it is for that army to travel along that road). The affinities of the roads for the various armies are fixed for the duration of the exercise. We also fix an integer MAX that is greater than or equal to all of the affinities.

The purpose of the exercise is to see how the final territories of each of the armies depend on their initial arrangements. Since we are discussing an algorithm here, no initiative is to be taken by the individual armies: they have to follow the rules of combat to be described momentarily.

All through the exercise each castle will have a *strength* assigned to it, this strength is an integer in the range $[0, \ldots, MAX]$. The strength of a castle may change as the exercise proceeds. Also, at any time, each castle may be occupied by one or more of the armies.

The exercise starts by distributing the soldiers of the armies into some of the castles, assigning to those castles that have soldiers in them the strength MAX, and to all other castles the strength 0. We say that this distribution of armies and strengths describes the situation at the start of Iteration 1.

The exercise proceeds in discrete iterative steps. The following gets done during Iteration k. Those soldiers (and only those soldiers) which occupy a castle of strength $MAX + 1 - k$ will try to increase the territory of their army. They will send units from their castle toward all the other castles. When these units arrive at another castle, their strength will be defined as the minimum of $MAX+1-k$ and the affinity for their army of the road from the originally occupied castle to the new one. If the strength of each of the armies arriving at a castle is less than the strength of the castle, then the castle's strength and occupancy will not change. If at least one arriving army has strength equal to that of the castle but no arriving army has greater strength, then the strength of the castle will not change, but it will get occupied also

[1] In fact, we prove this claim under slightly weaker hypotheses than those stated in [8].

by those arriving armies whose strength matches its strength (but not by any of the others). If some of the arriving armies have greater strength than the strength of the castle, then the castle will be taken over by those (and only those) arriving armies that have the greatest strength, and the strength of the castle will be set to the strength of the new occupiers. This describes what happens at iteration k except for one detail: if an army gets to occupy a new castle because its strength is $MAX + 1 - k$ (this can only happen if the affinity for this army of the road to this castle is at least $MAX + 1 - k$), then that army is allowed to send out units from this new castle as well.

The exercise stops at the end of Iteration *MAX*. The output of the algorithm provides, for each castle, the strength of the castle and the armies that occupy it at the end of the exercise.

2. Theory

In our very general approach we deal with an arbitrary finite set V, whose elements are referred to as *spels* (short for *spatial elements*). These spels can represent many different things, such as pixels of an image (as in [3,10,12,15]), dots in the plane (as in [16]) or feature vectors (as in [6]). In the picturesque description above, V is the set of castles. Furthermore, the theory and algorithm introduced in [8], and further discussed here, are independent of the specifics of the application area, and thus can be applied to data clustering [9] in general. A special choice (some papers on fuzzy segmentation restrict their attention only to V's of this type, see for example [14]) is when the V is of the form

$$V = \{c \in \mathbb{Z}^n \mid -b_j \leqslant c_j \leqslant b_j \text{ for some } b \in \mathbb{Z}^n_+\}, \tag{1}$$

where $\mathbb{Z}^n_+$ is the set of n-tuples of positive integers. Throughout this paper we illustrate the methods on a particularly simple V, which we denote by $\overline{V}$, which is defined by (1) with $n = 1$ and $b_1 = 1$ (i.e., $\overline{V} = \{(-1), (0), (1)\}$).

We desire to partition V into a number of objects, but in a fuzzy way; i.e., in addition to a spel being judged to belong to a particular object, it is also assigned a grade of membership in the object (that is, a number between 0 and 1, where 0 indicates that the spel definitely does not belong to the object, and 1 indicates that it definitely does). In the picturesque description above, at the end of the exercise each object consists of all the castles occupied by one particular army, and the grade of membership of the castle in the object is proportional to its strength. (To make the grade of membership satisfy the requirement that it is not greater than 1, we can divide the strength of each castle by *MAX*.) To formalize such fuzzy partitioning, we introduce the concept of an M-semisegmentation (where M is the number of objects).

An *M-semisegmentation* of V is a function σ that maps each $c \in V$ into an $(M + 1)$-dimensional vector $\sigma^c = (\sigma^c_0, \sigma^c_1, \ldots, \sigma^c_M)$, such that

1. $\sigma^c_0 \in [0, 1]$ (i.e., σ_0 is nonnegative but not greater than 1);
2. for each m $(1 \leqslant m \leqslant M)$, the value of σ^c_m is either 0 or σ^c_0; and
3. for at least one m $(1 \leqslant m \leqslant M)$, $\sigma^c_m = \sigma^c_0$.

Here σ_m^c represents the grade of membership of the spel c in the mth object, and σ_0^c is evidently always $\max_{1 \leqslant m \leqslant M} \sigma_m^c$.

We point out that this definition of M-semisegmentation allows a spel to belong to more than one object, as long as it has the same grade of membership in all of them. For the exercise described above, σ_0^c is proportional to the strength of the castle c and the mth army occupies that castle if, and only if, $\sigma_m^c = \sigma_0^c > 0$.

We say that an M-semisegmentation σ is an *M-segmentation* if, for every spel c, σ_0^c is positive. An example of a 2-segmentation $\bar{\sigma}$ of $\overline{V}$ is defined by $\bar{\sigma}^{(-1)} = (1, 0, 1)$, $\bar{\sigma}^{(0)} = (1, 1, 0)$ and $\bar{\sigma}^{(1)} = (0.25, 0.25, 0)$; i.e., (-1) is definitely in the second object, (0) is definitely in the first object, and (1) is in the first object with grade of membership 0.25.

The M-semisegmentations in our theory will be determined by M-fuzzy graphs, a concept that we now proceed to define.

We call a sequence $\langle c^{(0)}, \ldots, c^{(K)} \rangle$ of distinct spels a *chain*; its *links* are the ordered pairs $(c^{(k-1)}, c^{(k)})$ of consecutive spels in the sequence. The strength of a link is also a fuzzy concept (i.e., for every ordered pair (c, d) of spels, we assign a real number not less than 0 and not greater than 1, which we define as the strength of the link from c to d). To be precise, the *ψ-strength of a link* is provided by the appropriate value of a *fuzzy spel affinity* function $\psi : V^2 \to [0, 1]$, i.e., a function that assigns a value between 0 and 1 to every ordered pair of spels in V. (As we illustrate later, for the purpose of fuzzy segmentation of images, fuzzy spel affinities can often be automatically defined based on statistical properties of the links within regions identified by the user as belonging to the object of interest.) The *ψ-strength of a chain* is the ψ-strength of its weakest link; the ψ-strength of a chain with only one spel in it is 1 by definition. A set $U (\subseteq V)$ is said to be *ψ-connected* if, for every pair of spels in U, there is a chain in U of positive ψ-strength from the first spel of the pair to the second. For the picturesque description above, (c, d) denotes the one-way road from castle c to castle d, and an affinity of an army for this road has to be divided by *MAX* in order to match the definition of a fuzzy spel affinity.

In our approach there are no further restrictions on the definition of fuzzy spel affinity. Other researchers (e.g., [14]) restrict them to be reflexive (i.e., $\psi(c, c) = 1$ for all $c \in V$) and, much more significantly, to be symmetric (i.e., $\psi(c, d) = \psi(d, c)$ for all $c, d \in V$). Examples of such reflexive and symmetric fuzzy spel affinities are $\overline{\psi_1}$ and $\overline{\psi_2}$, defined by the additional conditions $\overline{\psi_1}((-1), (0)) = 0.5$, $\overline{\psi_1}((0), (1)) = 0.25$ and $\overline{\psi_1}((-1), (1)) = 0$, and $\overline{\psi_2}((-1), (0)) = \overline{\psi_2}((0), (1)) = 0.5$ and $\overline{\psi_2}((-1), (1)) = 0$. The $\overline{\psi_1}$-strength of the chain $\langle (-1), (0), (1) \rangle$ in $\overline{V}$ is 0.25, its $\overline{\psi_2}$-strength is 0.5, and $\overline{V}$ is both $\overline{\psi_1}$-connected and $\overline{\psi_2}$-connected.

If one wants to segment multiple objects, it is reasonable to define different fuzzy spel affinities for each one of them. (This corresponds to the idea of each army having its own affinity for each one-way road.) In general, an *M-fuzzy graph* is a pair (V, Ψ), where V is a nonempty finite set and $\Psi = (\psi_1, \ldots, \psi_M)$ with ψ_m (for $1 \leqslant m \leqslant M$) being a fuzzy spel affinity. An example of a 2-fuzzy graph is $(\overline{V}, \overline{\Psi})$, where $\overline{\Psi} = (\overline{\psi_1}, \overline{\psi_2})$. An M-fuzzy graph can be used to totally specify the aspects of the castles and the roads connecting them that are relevant to the rules of combat given above.

A *seeded* M-fuzzy graph is a triple $(V, \Psi, \mathscr{V})$ such that (V, Ψ) is an M-fuzzy graph and $\mathscr{V} = (V_1, \ldots, V_M)$, where $V_m \subseteq V$ for $1 \leqslant m \leqslant M$. A seeded M-fuzzy graph

$(V, (\psi_1, \ldots, \psi_M), (V_1, \ldots, V_M))$ is *connectable* if

1. the set V is ϕ_Ψ-connected, where $\phi_\Psi(c,\ d) = \min_{1 \leqslant m \leqslant M} \psi_m(c, d)$ for all $c,\ d \in V$, and
2. $V_m \neq \emptyset$, for at least one m, $1 \leqslant m \leqslant M$.

For an M-semisegmentation σ of V and for $1 \leqslant m \leqslant M$, we define the chain $\langle c^{(0)}, \ldots, c^{(K)} \rangle$ to be a *σm-chain* if $\sigma_m^{c^{(k)}} > 0$, for $0 \leqslant k \leqslant K$. Furthermore, for $W \subseteq V$ and $c \in V$, we use $\mu_{\sigma,m,W}(c)$ to denote the maximal ψ_m-strength of a σm-chain from a spel in W to c. (This is 0 if there is no such chain.)

Theorem 1. *If $(V, \Psi, \mathscr{V})$ is a seeded M-fuzzy graph (where $\Psi = (\psi_1, \ldots, \psi_M)$ and $\mathscr{V} = (V_1, \ldots, V_M)$), then*

(i) *there exists an M-semisegmentation σ of V with the following property: for every $c \in V$, if for $1 \leqslant n \leqslant M$*

$$s_n^c = \begin{cases} 1 & \text{if } c \in V_n, \\ \max\limits_{d \in V} (\min(\mu_{\sigma,n,V_n}(d), \psi_n(d, c))) & \text{otherwise}, \end{cases} \tag{2}$$

then for $1 \leqslant m \leqslant M$

$$\sigma_m^c = \begin{cases} s_m^c & \text{if } s_m^c \geqslant s_n^c, \text{ for } 1 \leqslant n \leqslant M, \\ 0 & \text{otherwise}; \end{cases} \tag{3}$$

(ii) *this M-semisegmentation is unique; and*

(iii) *it is an M-segmentation, provided that $(V, \Psi, \mathscr{V})$ is connectable.*

Before discussing the validity of Theorem 1, let us discuss in less mathematical terms what it says. The property stated in Theorem 1 is a reasonable one, as we can see in Fig. 1. Suppose, as in Fig. 1, that c is an arbitrary spel and that σ^d is known for all other spels d. Then, for $1 \leqslant n \leqslant M$ ($M = 3$ in Fig. 1), the s_n^c of (2) is the maximal ψ_n-strength of a chain $\langle d^{(0)}, \ldots, d^{(L)}, c \rangle$ from a seed spel in V_n to c such that $\sigma_n^{d^{(l)}} > 0$ (i.e., $d^{(l)}$ belongs to the nth object) for $0 \leqslant l \leqslant L$. ($s_n^c$ is defined to be 0 if there is no such chain.) Intuitively, the mth object (the red, green or blue object) can "claim" that c belongs to it if, and only if, s_m^c is maximal. This is indeed how things get sorted out in (3): σ_m^c has a positive value only for such objects. Furthermore, this property tells us how any one spel relates to the various objects, provided that we know σ for the other spels: for a fixed spel c we can work out the values of the s_n^c using (2) and what we request is that, at that spel c, (3) be satisfied. Theorem 1 says that there is one, and only one, M-semisegmentation which satisfies this reasonable property simultaneously everywhere, and that this M-semisegmentation is in fact an M-segmentation provided that the seeded M-fuzzy graph is connectable.

We now illustrate Theorem 1 for the above-specified 2-fuzzy graph $(\overline{V}, \overline{\Psi})$. If we choose the sets of seed spels to be $\overline{V_1} = \{(0)\}$ and $\overline{V_2} = \{(-1)\}$, then we get exactly the $\overline{\sigma}$ defined above. Suppose, for example, that we have been informed that $\overline{\sigma}^{(-1)} = (1, 0, 1)$ and $\overline{\sigma}^{(0)} = (1, 1, 0)$ and we wish to use Theorem 1 to determine $\overline{\sigma}^{(1)}$. We find that $s_1^{(1)} = 0.25$ (obtained by the

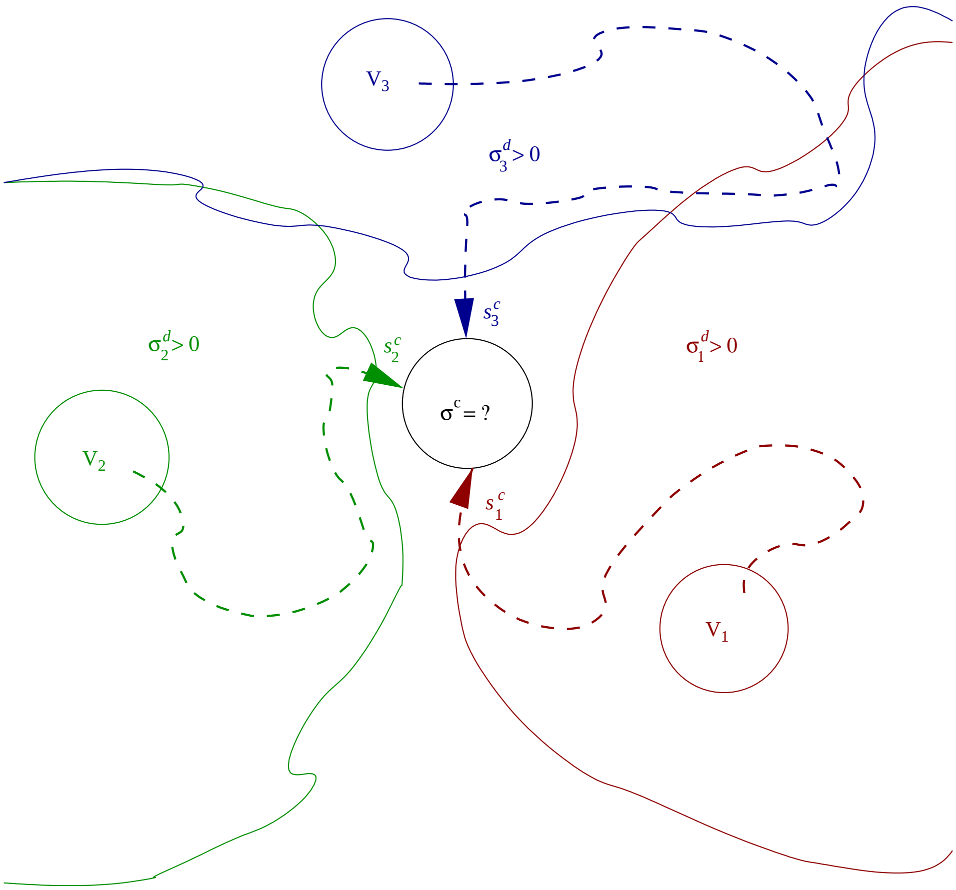

Fig. 1. Illustration of the desirability of the *M*-semisegmentation whose existence (and uniqueness) is guaranteed by Theorem 1. This figure appears in color in the online version of the paper.

choice $d=(0)$) and $s_2^{(1)}=0$ (if in (2) we choose d to be (-1), then $\overline{\psi_2}((-1),(1))=0$; if we choose it to be (0), then $\mu_{\overline{\sigma},2,\overline{V_1}}(0)=0$ since there is no $\overline{\sigma}$2-chain containing (0), due to the fact that $\overline{\sigma}_2^{(0)}=0$). Hence (3) tells us that indeed $\overline{\sigma}^{(1)}=(0.25,0.25,0)$. There is something subtle that takes place here: there is a chain $\langle(-1),(0),(1)\rangle$ of $\overline{\psi_2}$-strength 0.5 from the only seed spel of Object 2 to (1), while the maximal $\overline{\psi_1}$-strength of any chain from the only seed spel of Object 1 to (1) is only 0.25; nevertheless, (1) is assigned to Object 1 by Theorem 1, since the fact that (0) is a seed spel of Object 1 prevents it (for the given $\overline{\Psi}$) from being also in Object 2, and so the chain $\langle(-1),(0),(1)\rangle$ is "blocked" from being a $\overline{\sigma}$2-chain.

The proof of Theorem 1(i) shown below has not been published before (except in the preliminary version of this paper [4]), while the proofs of Theorem 1(ii) and Theorem 1(iii) were originally published in [8].

Proof of Theorem 1(i). In this existence proof we provide an inductive definition that resembles both the picturesque description of the previous section and the actual algorithm of the next section. The reader should, however, be warned: this inductive definition is not strictly identical to the algorithm (it was designed to make our proof simple, while the algorithm was designed to be efficient). In the next section we discuss the relationship between the inductive definition and the actual algorithm.

Let $R = \{1\} \cup \{\psi_m(c,d) > 0 \mid 1 \leqslant m \leqslant M, c, d \in V\}$. R is a finite set of real numbers from $(0, 1]$, and so its elements can be put into a strictly decreasing order $1 = {}^1r > {}^2r > \cdots > {}^{|R|}r > 0$. We will inductively define a sequence of M-semisegmentations ${}^1\sigma, {}^2\sigma, \ldots, {}^{|R|}\sigma$, and we will show that the M-semisegmentation ${}^{|R|}\sigma$ has the property stated in Theorem 1(i).

For any $c \in V$ and $1 \leqslant m \leqslant M$, we define

$$ {}^1\sigma_m^c = \begin{cases} 1 & \text{if there is a chain of } \psi_m\text{-strength 1 from a seed in } V_m \text{ to } c, \\ 0 & \text{otherwise.} \end{cases} \tag{4} $$

(Here, and later, the definition of ${}^i\sigma_0^c$ implicitly follows from the fact that ${}^i\sigma$ is an M-semisegmentation.)

For $1 \leqslant i \leqslant |R|$, we define

$$ {}^iU = \{c \in V \mid {}^i\sigma_0^c \geqslant {}^ir\}. \tag{5} $$

For $1 < i \leqslant |R|$, $c \in V$ and $1 \leqslant m \leqslant M$, we define

$$ {}^i\sigma_m^c = \begin{cases} {}^{(i-1)}\sigma_m^c & \text{if } c \in {}^{(i-1)}U, \\ {}^ir & \text{if there is a chain } \langle c^{(0)}, \ldots, c^{(K)} \rangle \text{ with } K \geqslant 1 \text{ of} \\ & \psi_m\text{-strength at least } {}^ir \text{ such that } c^{(0)} \in {}^{(i-1)}U, \\ & {}^{(i-1)}\sigma_m^{c^{(0)}} > 0,\ c^{(K)} = c \text{ and, for } 1 \leqslant k \leqslant K,\ c^{(k)} \notin {}^{(i-1)}U, \\ 0 & \text{otherwise.} \end{cases} \tag{6} $$

As an aid to understanding the implications of this definition, we mention the easily provable fact (not used in the proof that follows) that if the chain $\langle c^{(0)}, \ldots, c^{(K)} \rangle$ in (6) exists then the ψ_m-strength of the link $(c^{(0)}, c^{(1)})$, and hence the ψ_m-strength of the entire chain, must actually be exactly ${}^ir$.

It is obvious from these definitions that ${}^i\sigma$ is an M-semisegmentation, for $1 \leqslant i \leqslant |R|$. We now demonstrate the definitions on the already discussed 2-fuzzy graph $(\overline{V}, \overline{\Psi}, (\overline{V_1}, \overline{V_2}))$. For this case $R = \{1, 0.5, 0.25\}$. It immediately follows from (4) that ${}^1\sigma^{(-1)} = (1, 0, 1)$, ${}^1\sigma^{(0)} = (1, 1, 0)$, and ${}^1\sigma^{(1)} = (0, 0, 0)$. It turns out that ${}^2\sigma = {}^1\sigma$. This is because ${}^1U = \{(-1), (0)\}$, and there are no chains starting at either of these spels which satisfy all the conditions listed in the second line of (6). On the other hand, the chain $\langle (0), (1) \rangle$ can be used to generate ${}^3\sigma$,which is in fact the 2-segmentation specified by the condition of Theorem 1. This is not an accident, we are now going to prove that the ${}^{|R|}\sigma$ defined by (4)–(6) always has the property stated in Theorem 1(i).

It clearly follows from the definitions (4)–(6) that, for $c \in V$ and $1 \leqslant m \leqslant M$, ${}^{|R|}\sigma_m^c \in R \cup \{0\}$. Furthermore, it is also not difficult to see, for $1 \leqslant i \leqslant |R|$, that if $c \in {}^iU$, then ${}^i\sigma_m^c = {}^{|R|}\sigma_m^c$, and that

$$ {}^iU = \{c \in V \mid {}^{|R|}\sigma_0^c \geqslant {}^ir\}. \tag{7} $$

These facts imply the following two properties of the M-semisegmentation ${}^{|R|}\sigma$.

(**A**) For $c \in V$ and $1 \leqslant m \leqslant M$, ${}^{|R|}\sigma_m^c = 1$ if, and only if, there is a chain of ψ_m-strength 1 from a seed in V_m to c.
(**B**) For $c \in V$, $1 \leqslant m \leqslant M$ and $2 \leqslant i \leqslant |R|$, ${}^{|R|}\sigma_m^c = {}^i r$ if, and only if, there is a chain $\langle c^{(0)}, \ldots, c^{(K)} \rangle$ with $K \geqslant 1$ of ψ_m-strength at least ${}^i r$ such that $c^{(0)} \in {}^{(i-1)}U$, ${}^{|R|}\sigma_m^{c^{(0)}} > 0$, $c^{(K)} = c$ and, for $1 \leqslant k \leqslant K$, $c^{(k)} \notin {}^{(i-1)}U$.

Let $c, d \in V$. We say that (c, d) is *consistent* if either

$$c = d \tag{8}$$

or, for each m ($1 \leqslant m \leqslant M$), one of the following is true:

$${}^{|R|}\sigma_0^d > \min({}^{|R|}\sigma_m^c, \psi_m(c, d)); \tag{9}$$

$${}^{|R|}\sigma_0^d = \min({}^{|R|}\sigma_m^c, \psi_m(c, d)) = {}^{|R|}\sigma_m^d. \tag{10}$$

We now show that, for all $c, d \in V$, (c, d) is consistent.

To do this, we assume that there is a (c, d) and an m such that none of (8)–(10) holds, and show that this leads to a contradiction. A consequence of our assumption is that $c \neq d$ and that at least one of the following must be the case:

$${}^{|R|}\sigma_0^d < \min({}^{|R|}\sigma_m^c, \psi_m(c, d)); \tag{11}$$

$${}^{|R|}\sigma_0^d = \min({}^{|R|}\sigma_m^c, \psi_m(c, d)) \quad \text{and} \quad {}^{|R|}\sigma_m^d \neq {}^{|R|}\sigma_0^d. \tag{12}$$

We may assume that ${}^{|R|}\sigma_m^c > 0$ and that $\psi_m(c, d) > 0$, for otherwise one of (9) or (10) clearly holds. Hence ${}^{|R|}\sigma_m^c = {}^{|R|}\sigma_0^c = {}^i r$, for some $1 \leqslant i \leqslant |R|$. From (11) or (12) it follows that ${}^{|R|}\sigma_0^d \leqslant {}^i r$. Consequently, (7) implies that if $i \geqslant 2$, then neither c nor d is in ${}^{(i-1)}U$.

If $i = 1$, then by **A** there is a chain of ψ_m-strength 1 from a seed in V_m to c. If $i \geqslant 2$, then by **B** there is a chain $\langle c^{(0)}, \ldots, c^{(K)} \rangle$ with $K \geqslant 1$ of ψ_m-strength at least ${}^i r$ such that $c^{(0)} \in {}^{(i-1)}U$, ${}^{|R|}\sigma_m^{c^{(0)}} > 0$, $c^{(K)} = c$ and, for $1 \leqslant k \leqslant K$, $c^{(k)} \notin {}^{(i-1)}U$. Suppose, for now, that $\psi_m(c, d) \geqslant {}^i r$. In both of the just mentioned cases ($i = 1$ and $i \geqslant 2$), ${}^{|R|}\sigma_m^d = {}^{|R|}\sigma_0^d = {}^i r$. (This can be seen by observing that d is either already in the chain to c or the chain to c can be extended to d by the additional link (c, d).) But then (10) holds, a contradiction. So assume that $\psi_m(c, d) = {}^j r$ for some $j > i$. Since (11) or (12) holds, we get from (7) that $d \notin {}^{(j-1)}U$. But $c \in {}^{(j-1)}U$, and so, applying **B** to the chain $\langle c, d \rangle$, we get that ${}^{|R|}\sigma_m^d = {}^j r$. This implies that (10) holds. This final contradiction completes our proof that, for all $c, d \in V$, (c, d) is consistent.

Next we show that, for all $c \in V$ and $1 \leqslant m \leqslant M$,

$${}^{|R|}\sigma_m^c = \mu_{{}^{|R|}\sigma, m, V_m}(c). \tag{13}$$

To simplify the notation, we use s in this proof to abbreviate ${}^{|R|}\sigma_m^c$. Recall that $\mu_{{}^{|R|}\sigma, m, V_m}(c)$ denotes the maximal ψ_m-strength of an ${}^{|R|}\sigma m$-chain from a seed in V_m to c. Note that we can assume that $s \in R$, for the alternative is that $s = 0$ in which case there can be no ${}^{|R|}\sigma m$-chain that includes c and so that right-hand side of (13) is also 0 by definition. Our proof will be in two stages: first we show that there is an ${}^{|R|}\sigma m$-chain from a seed in V_m to c of ψ_m-strength at least s and then we show that there is no ${}^{|R|}\sigma m$-chain from a seed in V_m to c of ψ_m-strength greater than s.

To show the existence of an ${}^{|R|}\sigma m$-chain from a seed in V_m to c of ψ_m-strength at least s, we use an inductive argument. If $s = {}^1r = 1$, then the desired result is assured by **A**. Moreover (also by **A**), in this case the ${}^{|R|}\sigma$-chain with the stated properties lies in ${}^1U$. Now let $i > 1$ and $s = {}^ir$. Assume as induction hypothesis that, for $1 \leqslant j < i$, whenever a spel d is such that ${}^{|R|}\sigma_m^d = {}^jr$, there is an ${}^{|R|}\sigma m$-chain in ${}^jU$ from a seed in V_m to d of ψ_m-strength at least ${}^jr$.

By **B** there is a chain $\langle c^{(0)}, \ldots, c^{(K)}\rangle$ with $K \geqslant 1$ of ψ_m-strength at least s such that $c^{(0)} \in {}^{(i-1)}U$, ${}^{|R|}\sigma_m^{c^{(0)}} > 0$, $c^{(K)} = c$ and, for $1 \leqslant k \leqslant K$, $c^{(k)} \notin {}^{(i-1)}U$. We are now going to show that $\langle c^{(0)}, \ldots, c^{(K)}\rangle$ is an ${}^{|R|}\sigma m$-chain in ${}^iU$ by showing that, for $1 \leqslant k \leqslant K$, ${}^{|R|}\sigma_m^{c^{(k)}} = s$. Otherwise, consider the smallest $k \geqslant 1$ that violates this equation. Then we have that ${}^{|R|}\sigma_m^{c^{(k-1)}} \geqslant s$ and ${}^{|R|}\sigma_m^{c^{(k)}} \neq s$, but ${}^{|R|}\sigma_0^{c^{(k)}} \leqslant s$ (since $c^{(k)} \notin {}^{(i-1)}U$). This combined with the fact that $\psi_m(c^{(k-1)}, c^{(k)}) \geqslant s$ violates the consistency of $(c^{(k-1)}, c^{(k)})$. Since $c^{(0)} \in {}^{(i-1)}U$ and ${}^{|R|}\sigma_m^{c^{(0)}} > 0$, ${}^{|R|}\sigma_m^{c^{(0)}} = {}^jr$ for some $1 \leqslant j < i$ and so, by the induction hypothesis, there is an ${}^{|R|}\sigma m$-chain in ${}^jU$ from a seed in V_m to $c^{(0)}$ of ψ_m-strength at least ${}^jr > s$. Appending $\langle c^{(1)}, \ldots, c^{(K)}\rangle$to this chain we obtain an ${}^{|R|}\sigma m$-chain in ${}^iU$ from a seed in V_m to c of ψ_m-strength at least ${}^ir = s$. (No spels occur more than once in the resulting sequence since $c^{(1)}, \ldots, c^{(K)}$ do not belong to ${}^{(i-1)}U$, while the other elements of this sequence belong to ${}^jU$ for a $j < i$.)

Now we show that there is no ${}^{|R|}\sigma m$-chain from a seed in V_m to c of ψ_m-strength greater than s. This is clearly so if $s = 1$. Suppose now that $s < 1$ and that $\langle c^{(0)}, \ldots, c^{(K)}\rangle$ is an ${}^{|R|}\sigma m$-chain from a seed in V_m of ψ_m-strength $t > s$. We now show that, for $0 \leqslant k \leqslant K$, ${}^{|R|}\sigma_m^{c^{(k)}} \geqslant t$. From this it follows that $c^{(K)}$ cannot be c and we are done. Since $c^{(0)}$ is a seed in V_m, ${}^{|R|}\sigma_m^{c^{(0)}} = 1 \geqslant t$. For $k > 0$, the induction that makes use of the consistency of $(c^{(k-1)}, c^{(k)})$ leads to the desired result.

To show that $\sigma = {}^{|R|}\sigma$ satisfies the property stated in Theorem 1(i), we first make two preliminary observations:

(**a**) For any $c \in V$ and $1 \leqslant n \leqslant M$, if $\sigma_n^c > 0$, then $s_n^c = \sigma_n^c = \sigma_0^c$. (The first equality follows from (2) and (13), and the second from the definition of an M-semisegmentation.)

(**b**) For any $c \in V$ and $1 \leqslant n \leqslant M$, if $\sigma_n^c = 0$ and $\sigma_0^c > 0$, then $s_n^c < \sigma_0^c$. (Assume the contrary. It cannot be that s_n^c is defined by the first line of (2), for then $c \in V_n$ and by **A** we would have that $\sigma_n^c = 1$. Hence s_n^c is defined by the second line of (2) using some d such that $\min(\mu_{\sigma,n,V_n}(d), \psi_n(d, c)) = s_n^c \geqslant \sigma_0^c > 0$. Hence, by (13), $\sigma_n^d \geqslant \sigma_0^c > 0$. Interchanging c and d in the definition of consistency, we see that (8) cannot hold (since $\sigma_n^d > 0$ and $\sigma_n^c = 0$), (9) cannot hold (since $\sigma_0^c \leqslant \sigma_n^d$ and $\sigma_0^c \leqslant \psi_n(d, c)$), and (10) cannot hold (since $\sigma_n^c = 0$ and $\sigma_0^c > 0$). This contradiction with the consistency of (d, c) proves **b**.)

To complete the proof, let $c \in V$. We first assume that $\sigma_0^c = 0$. Then, by the definition of an M-semisegmentation, $\sigma_n^c = 0$ for $1 \leqslant n \leqslant M$. It follows from **A** that $c \notin V_n$ and so s_n^c is defined by the second line of (2). Thus if s_n^c were greater than 0, then there would have to be a $d \in V$ such that $\min(\mu_{\sigma,n,V_n}(d), \psi_n(d,c)) > 0$. Here $\mu_{\sigma,n,V_n}(d) = \sigma_n^d$ by (13) and so we also have $\min(\sigma_n^d, \psi_n(d,c)) > 0$. This and the consistency of (d,c) together imply that $\sigma_0^c > 0$, contrary to our assumption. Hence $s_n^c = 0$, and since this is true for $1 \leqslant n \leqslant M$, (3) holds for $1 \leqslant m \leqslant M$.

We now assume that $\sigma_0^c > 0$. By the definition of an M-semisegmentation, for $1 \leqslant n \leqslant M$, either $\sigma_n^c = \sigma_0^c$ (and there is at least one such n) or $\sigma_n^c = 0$. In the first case, we have by **a** that $s_n^c = \sigma_n^c = \sigma_0^c$, and in the second case, we have by **b** that $s_n^c < \sigma_0^c$. It again follows that (3) holds for $1 \leqslant m \leqslant M$. $\square$

Proof of Theorem 1(ii). Suppose that there are two different M-semisegmentations σ and τ of V having the stated property. We choose a spel c, such that $\sigma^c \neq \tau^c$, but for all $d \in V$ such that $\max(\sigma_0^d, \tau_0^d) > \max(\sigma_0^c, \tau_0^c)$, $\sigma^d = \tau^d$. Without loss of generality, we assume that $\sigma_0^c \geqslant \tau_0^c$, from which it follows that, for some $m \in \{1, \ldots, M\}$, $\sigma_m^c > \tau_m^c (\geqslant 0)$ and so, by (3), $\sigma_m^c = s_m^c$ and $c \notin V_m$. This implies that there exists a σm-chain $\langle d^{(0)}, \ldots, d^{(L)} \rangle$ in V of ψ_m-strength not less than $\sigma_m^c (> 0)$ such that $d^{(0)} \in V_m$ and $\psi_m(d^{(L)}, c) \geqslant \sigma_m^c$. Next we show that $\langle d^{(0)}, \ldots, d^{(L)} \rangle$ is a τm-chain.

We need to show that, for $0 \leqslant l \leqslant L$, $\tau_m^{d^{(l)}} > 0$. This is true for 0, since $d^{(0)} \in V_m$. Now assume that it is true for $l-1$ $(1 \leqslant l \leqslant L)$. Since $\langle d^{(0)}, \ldots, d^{(l-1)} \rangle$ is a τm-chain in V of ψ_m-strength at least $\sigma_m^c (> 0)$ from an element of V_m, we have that $\mu_{\tau,m,V_m}(d^{(l-1)}) \geqslant \sigma_m^c$. Since we also know that $\psi_m(d^{(l-1)}, d^{(l)}) \geqslant \sigma_m^c$, we get that $t_m^{d^{(l)}} \geqslant \sigma_m^c$ (where t is defined for τ as s is defined for σ in (2)). The only way $\tau_m^{d^{(l)}}$ could be 0 is if there were an $n \in \{1, \ldots, M\}$ such that $t_n^{d^{(l)}} > t_m^{d^{(l)}}$. Then $\max(\sigma_0^{d^{(l)}}, \tau_0^{d^{(l)}}) \geqslant \tau_0^{d^{(l)}} = \tau_n^{d^{(l)}} = t_n^{d^{(l)}} > t_m^{d^{(l)}} \geqslant \sigma_m^c = \sigma_0^c = \max(\sigma_0^c, \tau_0^c)$. By the choice of c, this would imply that $\sigma^{d^{(l)}} = \tau^{d^{(l)}}$, which cannot be since $\sigma_m^{d^{(l)}} \neq 0$.

From the facts that $\langle d^{(0)}, \ldots, d^{(L)} \rangle$ is a τm-chain of ψ_m-strength not less than σ_m^c and that $\psi_m(d^{(L)}, c) \geqslant \sigma_m^c$, it follows that $\tau_0^c \geqslant t_m^c \geqslant \sigma_m^c = \sigma_0^c \geqslant \tau_0^c$, implying that all the inequalities are in fact equalities. But then $\sigma_m^c = t_m^c = \tau_m^c$, contradicting $\sigma_m^c > \tau_m^c$ and thereby validating uniqueness. $\square$

Proof of Theorem 1(iii). We observe that it is a consequence of (3) that, for any spel c, $\sigma_0^c = \max_{1 \leqslant m \leqslant M} s_m^c$. Let $\langle c^{(0)}, \ldots, c^{(K)} \rangle$ be a chain of positive ϕ_Ψ-strength from a seed spel to an arbitrary spel c. (Such a chain exists since $(V, \Psi, \mathscr{V})$ is assumed to be connectable.) We now show inductively that, for $0 \leqslant k \leqslant K$, $\sigma_0^{c^{(k)}} > 0$. This is clearly so for $k=0$. Suppose now that it is so for $k-1$. Choose an m $(1 \leqslant m \leqslant M)$ such that $\sigma_0^{c^{(k-1)}} = \sigma_m^{c^{(k-1)}} = s_m^{c^{(k-1)}}$. Then there is a σm-chain of positive ψ_m-strength from a spel in V_m to $c^{(k-1)}$. Since $\psi_m(c^{(k-1)}, c^{(k)}) > 0$, $\sigma_0^{c^{(k)}} \geqslant s_m^{c^{(k)}} > 0$. $\square$

Note that the proof of Theorem 1(i) gives us an alternative characterization of the unique M-semisegmentation that is determined by the property in Theorem 1(i). This is because in the proof we show that ${}^{|R|}\sigma$ is such that every pair of spels is consistent, as defined by (8)–(10), and that (13) is satisfied. In the rest of the proof it is only these facts and **A** that are used to show that ${}^{|R|}\sigma$ satisfies the property of Theorem 1(i). Hence our M-semisegmentation can also be uniquely characterized as one which satisfies **A** and (13) and for which every pair of spels is consistent.

3. Algorithm

We claim that the picturesque algorithm described in Section 1 produces an output that is essentially the M-segmentation σ of Theorem 1. However, a direct implementation of that algorithm would not be computationally efficient: many of the iterative steps would result in no change of the status quo, and even if changes were to take place during an iterative step, resources would be wasted on performing actions that can be avoided by a more carefully designed algorithm that aims at producing the same output.

In [8] the first two authors presented the efficient greedy MOFS (multi-object fuzzy segmentation) algorithm for this purpose; below we give a detailed specification of it. It makes use of a priority queue H (a binary heap) of spels c, with associated keys σ_0^c [5]. Such a priority queue has the property that the key of the spel at its head is maximal (its value is denoted by Maximum-Key(H), which is defined to be 0 if H is empty). As the algorithm proceeds, each spel is inserted into H exactly once (using the operation $H \leftarrow H \cup \{c\}$) and is eventually removed from H (using the operation Remove-Max(H), which removes the spel at the head of the priority queue). At the time when a spel c is removed from H, the vector σ^c has its final value. Spels are removed from H in a non-increasing order of the final value of σ_0^c. We use the variable r to store the current value of Maximum-Key(H).

MOFS algorithm

for $c \in V$ **do**
$\quad$ **for** $m \leftarrow 0$ **to** M **do**
$\quad\quad \sigma_m^c \leftarrow 0$
$H \leftarrow \emptyset$
for $m \leftarrow 1$ **to** M **do**
$\quad U_m \leftarrow V_m$
$\quad$ **for** $c \in U_m$ **do**
$\quad\quad$ **if** $\sigma_0^c = 0$ **then do** $H \leftarrow H \cup \{c\}$
$\quad\quad \sigma_0^c \leftarrow \sigma_m^c \leftarrow 1$
$r \leftarrow 1$
while $r > 0$ **do**
$\quad$ **for** $m \leftarrow 1$ **to** M **do**
$\quad\quad$ **while** $U_m \neq \emptyset$ **do**
$\quad\quad\quad$ remove a spel d from U_m
$\quad\quad\quad C \leftarrow \{c \in V \mid \sigma_m^c < \min(r, \psi_m(d,c))$ **and** $\sigma_0^c \leqslant \min(r, \psi_m(d,c))\}$
$\quad\quad\quad$ **while** $C \neq \emptyset$ **do**
$\quad\quad\quad\quad$ remove a spel c from C
$\quad\quad\quad\quad t \leftarrow \min(r, \psi_m(d,c))$
$\quad\quad\quad\quad$ **if** $r = t$ **and** $\sigma_m^c < r$ **then do** $U_m \leftarrow U_m \cup \{c\}$
$\quad\quad\quad\quad$ **if** $\sigma_0^c < t$ **then do**
$\quad\quad\quad\quad\quad$ **if** $\sigma_0^c = 0$ **then do** $H \leftarrow H \cup \{c\}$
$\quad\quad\quad\quad\quad$ **for** $n \leftarrow 1$ **to** M **do**
$\quad\quad\quad\quad\quad\quad \sigma_n^c \leftarrow 0$
$\quad\quad\quad\quad \sigma_0^c \leftarrow \sigma_m^c \leftarrow t$

25. **while** Maximum-Key$(H) = r$ **do**
26. Remove-Max(H)
27. $r \leftarrow$ Maximum-Key(H)
28. **for** $m \leftarrow 1$ **to** M **do**
29. $U_m \leftarrow \{c \in H \mid \sigma_m^c = r\}$

We now demonstrate the correctness of this algorithm in the sense that we indicate why it produces the ${}^{|R|}\sigma$ defined by (4)–(6). We do not consider it necessary to give a formal proof here; a discussion of the relationship of the operation of the MOFS algorithm to the definition should suffice.

The process is initialized (Steps 1–10) by first setting σ_m^c to 0, for each spel c and $0 \leqslant m \leqslant M$. Then, for every seed spel $c \in V_m$, c is put into U_m and into H and both σ_0^c and σ_m^c are set to 1. Following this, r is also set to 1. At the end of the initialization, the following conditions are satisfied. (We are assuming here that $V_m \neq \emptyset$ for at least one m. It is trivial to prove that the algorithm performs correctively in the alternative case.)

(i) σ is an M-semisegmentation of V.
(ii) A spel c is in H if, and only if, $0 < \sigma_0^c \leqslant r$.
(iii) $r =$ Maximum-Key(H).
(iv) For $1 \leqslant m \leqslant M$, $U_m = \{c \in H \mid \sigma_m^c = r\}$.

It would be nice for easy understanding of the relationship between the algorithm and the definition if σ at this stage were the same as the ${}^1\sigma$ of (4). However, this is not so: in (4) we assign value 1 not only to things in V_m, but also to things that can be reached from V_m by chains of ψ_m-strength 1. It is computationally more efficient to postpone and intermix this action with the next stage. Step 19 of the algorithm is what takes care of this, in a manner that we discuss momentarily.

The initialization is followed by the main loop of the algorithm. At the beginning of each execution of this loop, conditions (i)–(iv) above are satisfied. The main loop is repeatedly performed for decreasing values of r until r becomes 0, at which time the algorithm terminates (Step 11). There are two parts to the main loop, each of which has a very different function.

The first part of the main loop (Steps 12–24) is the essential part of the MOFS algorithm. It is in here where we update our best guess so far of the final values of the σ_m^c. A current value is replaced by a larger one if it is found that there is a σm-chain from a seed spel in V_m to c of ψ_m-strength greater than the old value (the previously maximal ψ_m-strength of the known σm-chains of this kind) and it is replaced by 0 if it is found that (for an $n \neq m$) there is a σn-chain from a seed spel in V_n to c of ψ_n-strength greater than the old value of σ_m^c.

To understand the relationship of the main loop of the algorithm to the definition in (6) consider the following. The r in the algorithm corresponds to the ${}^i r$ in the definition. When the loop is entered, the set U_m contains some (but not necessarily all) spels $c \in {}^i U$ for which ${}^i\sigma_m^c > 0$. However, as the execution of the loop proceeds, all spels that satisfy this condition will get put into U_m (in Step 19).

For the sake of computational efficiency, the algorithm does something that is not directly reflected in definition (6): as soon as an opportunity arises, it greedily estimates values ${}^j\sigma_m^c$

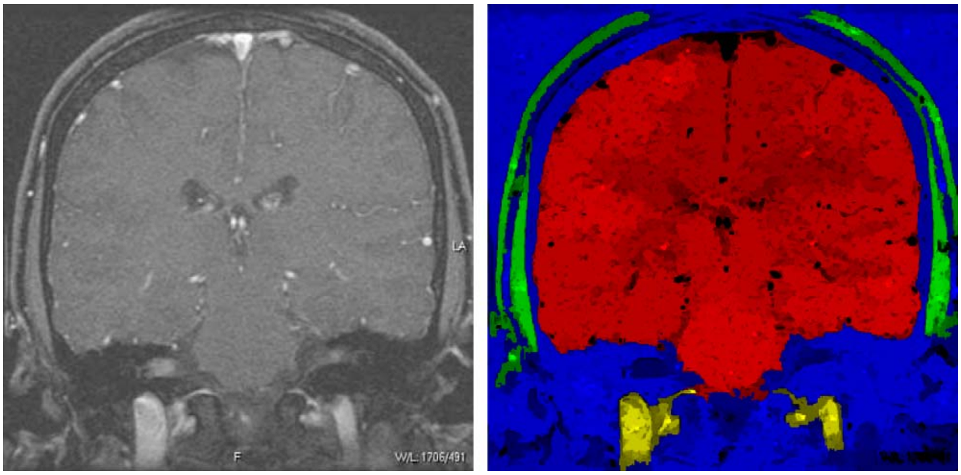

Fig. 2. MRI of a patient and a 4-segmentation of it. This figure appears in color in the online version of the paper.

for $j \geqslant i$. Although some of this effort may be wasted, in the sense that the estimated value will be replaced by another one later on, the greedy strategy allows us to avoid having to search explicitly for spels that satisfy the rather complicated condition in the second line of (6).

The purpose of the second part of the main loop (Steps 25–29) is to restore the satisfaction of conditions (iii) and (iv) above for a new (smaller) value of r. It is here that the use of the priority queue structure of H comes into its own: it allows us to skip over steps implied by the inductive definition during which nothing would happen (because we would have ${}^{i}\sigma = {}^{(i-1)}\sigma$).

4. Experiments

For our first illustration of the use of the MOFS algorithm we segmented a two-dimensional (2D) image defined on a V of the type specified in (1).[2] Fig. 2 shows a 400×397 magnetic resonance image (MRI) of a head on the left and a 4-segmentation of it on the right.

The way we specify ψ_m and V_m $(1 \leqslant m \leqslant 4)$ for such an image is the following. We click on some spels in the image to identify them as belonging to the mth object, and V_m is formed by these points and their eight neighbors. We define g_m to be the mean and h_m to be the standard deviation of the average brightness for all edge-adjacent pairs of spels in V_m and a_m to be the mean and b_m to be the standard deviation of the absolute differences of brightness for all edge-adjacent pairs of spels in V_m. Then we define $\psi_m(c, d)$ to be 0 if c and d are not edge-adjacent and to be $[\rho_{g_m,h_m}(g) + \rho_{a_m,b_m}(a)]/2$ if they are, where g is the mean and a is the absolute difference of the brightnesses of c and d and the function $\rho_{r,s}(x)$

[2] For examples using images defined on a hexagonal grid see [8].

is the probability density function of the Gaussian distribution with mean r and standard deviation s multiplied by a constant so that the peak value becomes 1.

For this segmentation we selected seed points belonging to various anatomically relevant parts (for example, the red seed points were used to identify brain tissue). The segmentation shown on the right of Fig. 2 actually tells us more than just to which object a spel belongs (as indicated by its hue), it also encodes in the brightness of each spel its grade of membership. In fact, one can identify the ventricular cavities inside the brain due to their having low brightness values in the red object.

The execution time needed by our implementation of the MOFS algorithm to segment the image shown in Fig. 2 was 1.26 s using an Intel® Xeon™ 1.7 GHz personal computer, or approximately 8 μs per spel.

As pointed out earlier, the multiseeded segmentation algorithm is general enough to be applied to images defined on various grids. We now illustrate the MOFS algorithm by segmenting a three-dimensional (3D) image defined on the face-centered cubic (fcc) grid. (Reasons for using such a grid are discussed in [7], especially in Chapter 2.)

Using $\mathbb{Z}$ for the set of all integers and δ for a positive real number, we define the *face-centered cubic* (fcc) grid F_δ by

$$F_\delta = \{(\delta c_1, \delta c_2, \delta c_3) \mid c_1, c_2, c_3 \in \mathbb{Z} \text{ and } c_1 + c_2 + c_3 \equiv 0 (\mathrm{mod}\ 2)\}, \tag{14}$$

where δ denotes the grid spacing. We define the adjacency relation β for the grid F_δ by: for any pair (c, d) of grid points in F_δ,

$$(c, d) \in \beta \Leftrightarrow \|c - d\| = \sqrt{2}\delta. \tag{15}$$

Each grid point $c \in F_\delta$ has 12 β-adjacent grid points in F_δ.

Experiments with segmentations using this approach on 3D images were reported in [1,2]. Here we show the results of one of the experiments reported in [1] that was performed on a Computerized Tomography (CT) reconstruction that assigned values to a total of $(298 \times 298 \times 164)/2 = 7,281,928$ (see (14)) fcc grid points. We selected seeds for four objects: the intestine (red object), other soft tissues (green object), the bones (blue object) and the lungs/background (cyan object). The corresponding fuzzy spel affinities were calculated in a manner strictly analogous to the above-discussed 2D example. Then, using a 1.7 GHz Intel® Xeon™ personal computer, our program performed the 4-segmentation on this volume that is illustrated in Fig. 3.

The execution time of our program was 249 s, or approximately 34 μs were needed per spel to perform the segmentation. Based on the execution time for the 2D experiments, one would expect a smaller execution time for the 3D volume. There are three main reasons why the average execution time needed per spel is higher. First, since we used β-adjacency, the number of neighboring spels was tripled as compared to the 2D example, where we used edge-adjacency (four neighbors). Second, a memory saving approach that we used in implementing the 3D version of our algorithm slowed down the execution. Finally, our program was developed with the goal of being able to segment images placed on a variety of 3D grids, and this generality also contributed to the longer execution time of the algorithm, as opposed to the approach taken in the 2D case, where we used a special purpose program to produce segmentations of images defined on grids of the type defined in (1).

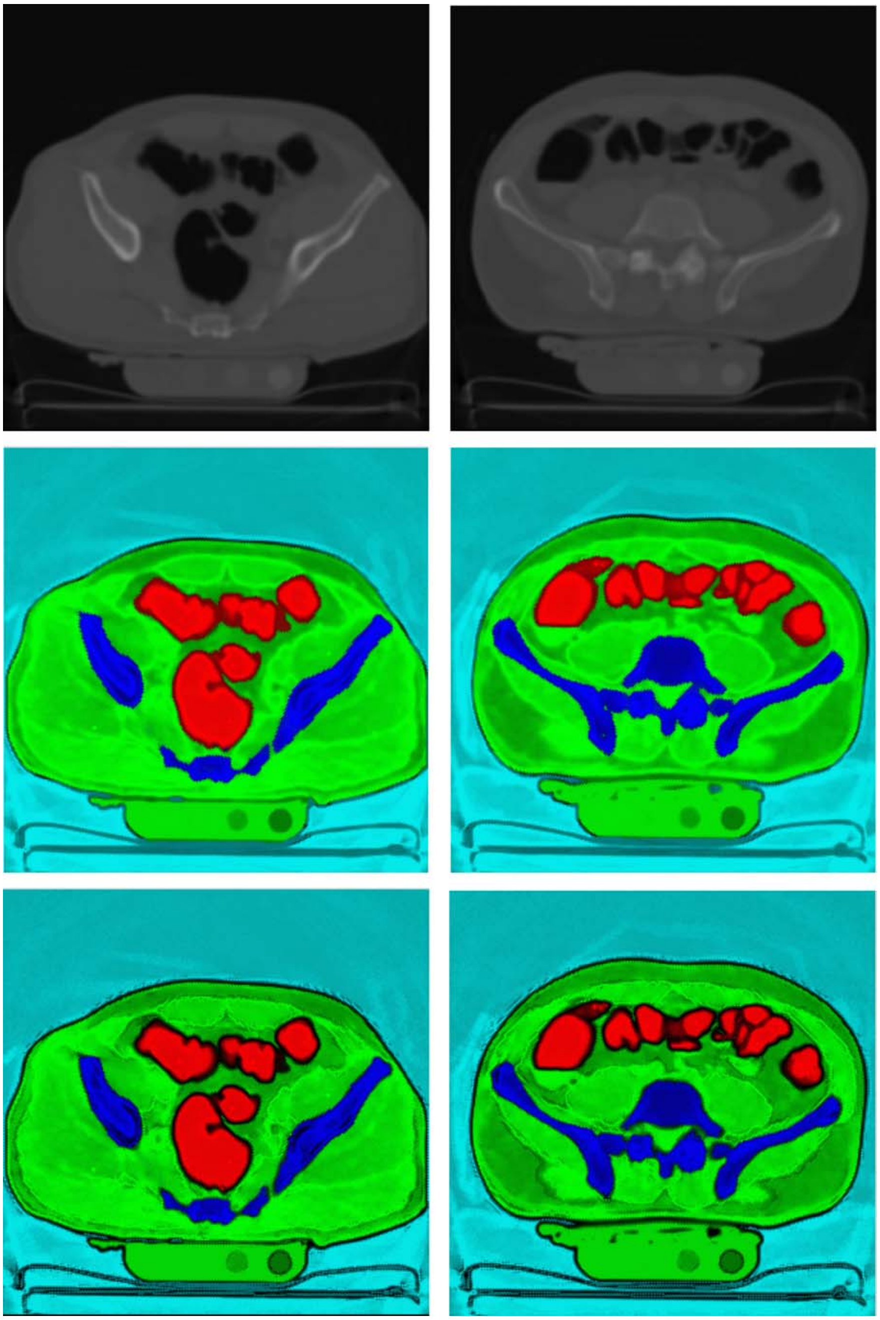

Fig. 3. Two axial slices from a CT volume placed on the fcc grid and the corresponding slices of 4-segmentations obtained using the original implementation of the MOFS algorithm (middle row) and the implementation discussed in Section 5 (bottom row). (All six images were interpolated for display purposes. This figure appears in color in the online version of the paper.)

5. Speeding up the algorithm

Even though we are able to segment a 3D image with more than 7,000,000 spels in approximately 4 min, this response time may not be reasonable for some applications. This prompted us to look into ways of speeding up the execution of the MOFS algorithm. Here we present a fast implementation of the MOFS algorithm that can be employed in certain circumstances.

Suppose that the set R of nonzero fuzzy spel affinities for a particular class of problems is always a subset of a fixed set A. Let K be the cardinality of the set $A \cup \{1\}$, and let $1 = a_1 > a_2 > \cdots > a_K > 0$ be the elements of A. For example, in many applications the quality of the fuzzy segmentation is not significantly affected if we round each fuzzy spel affinity to three decimal places. If we use such rounded spel affinities, then we can take $A = \{0.001, 0.002, \ldots, 0.999, 1.000\}$, so that $K = 1000$ and $a_k = 1.001 - k/1000$.

Our new implementation is presented below in pseudo-code. Instead of the priority queue H that was used in the first implementation, the new implementation uses an $M \times K$ array $U[m][k]$ of sets of nodes that represent spels, where M is (as before) the number of objects. (Similar ideas were used in [11] to speed up the algorithm of [15].) This implementation is most effective if all of its data structures (with space complexity $\mathrm{O}(M(K+V))$) can be held in the main memory.

Fast implementation of the MOFS algorithm

for $c \in V$ **do**
    for $m \leftarrow 0$ **to** M **do**
        $\sigma_m^c \leftarrow 0$
for $m \leftarrow 1$ **to** M **do**
    for $c \in V_m$ **do**
        $\sigma_0^c \leftarrow \sigma_m^c \leftarrow 1$
    $U[m][1] \leftarrow V_m$
    for $k \leftarrow 2$ **to** K **do**
        $U[m][k] \leftarrow \emptyset$
for $k \leftarrow 1$ **to** K **do**
    for $m \leftarrow 1$ **to** M **do**
        while $U[m][k] \neq \emptyset$ **do**
            remove a spel d from the set $U[m][k]$
            $C \leftarrow \{c \in V \mid \sigma_m^c < \min(a_k, \psi_m(d,c))$ **and** $\sigma_0^c \leqslant \min(a_k, \psi_m(d,c))\}$
            while $C \neq \emptyset$ **do**
                remove a spel c from C
                $t \leftarrow \min(a_k, \psi_m(d,c))$
                if $\sigma_0^c < t$ **then do**
                    remove c from each set in U that contains it

20. **for** $n \leftarrow 1$ **to** M **do**
21. $\sigma_n^c \leftarrow 0$
22. $\sigma_0^c \leftarrow \sigma_m^c \leftarrow t$
23. insert c into the set $U[m][l]$ where l is the integer such that $a_l = t$

The following informal discussion of the new implementation, in the language of the picturesque description given in the Introduction, may be helpful. At the start of Iteration k, for every index l in the range $k \leqslant l \leqslant MAX$, the set $U[m][l]$ holds all those castles that are occupied by army m and whose current strength is $MAX + 1 - l$. (In particular, at the start of Iteration k the set $U[m][k]$ consists of all the castles occupied by army m that have strength $MAX + 1 - k$.) During Iteration k, units of army m are sent from each castle in $U[m][k]$ to all other castles—see Steps 11–14 of the pseudo-code. When such a unit of army m succeeds in occupying another castle c, and the resulting strength of c is $MAX + 1 - l$, we insert c into the set $U[m][l]$—see Step 23 of the pseudo-code. But if c is taken over by such a unit (i.e., if c's previous strength was less than $MAX + 1 - l$), then we first remove c from any sets that contain c—see Step 19 of the pseudo-code.

Each set $U[m][k]$ is represented by a doubly linked list of nodes. Nodes in different sets that represent the same spel are chained together in a linked list, which makes it easy to remove a spel from all the sets that contain it. (There can be at most M nodes in such a linked list, since each $U[m][k]$ list contains at most one node that represents any given spel c, and all the $U[m][k]$ lists that contain nodes representing c must have the same index k.) A list node contains its spel's coordinates while the data of spel c includes a pointer to the linked list of nodes that represent c.

In the earlier implementation, the heap H must be updated at a time cost of $\mathrm{O}(\log N)$ each time the strength σ_0^c of a spel c increases, where N is the number of spels in the heap at that time. The final removal of a spel c from the heap also has a time cost of $\mathrm{O}(\log N)$. In the new implementation, when the value of any component σ_m^c changes we make corresponding changes in our sets of nodes, but the time cost of making those changes is $\mathrm{O}(\max(1, |\{m \geqslant 1 \mid \sigma_m^c > 0\}|)$. The cardinality of the set $\{m \geqslant 1 \mid \sigma_m^c > 0\}$ is bounded by M (the number of objects to be segmented), which is a small number (typically less than 10) in our current applications.

Fig. 3 shows two slices of the CT image mentioned at the end of Section 4 and the corresponding slices of two 4-segmentations computed using the original implementation (middle row) and the new implementation (bottom row) of the MOFS algorithm. To produce the latter 4-segmentation, the fuzzy affinity values were rounded to the third decimal place. Although this changed the M-fuzzy graph (V, Ψ) and, consequently, the M-semisegmentation whose existence and uniqueness are guaranteed by Theorem 1, the two 4-segmentations are in fact very similar. But the new implementation needed only 35 s of execution time, or approximately 5 μs per spel, to produce its 4-segmentation, compared with 249 s of execution time and approximately 34 μs per spel for the original implementation; i.e., we observed a speedup factor of 7.1. Fig. 4 shows one slice from an electron microscopy (EM) volume (top) and one slice each from two CT volumes, and the corresponding segmented slices obtained using the new implementation. The volumes contained 3,538,944 (top), 7,864,320 (middle) and 8,388,608 (bottom) spels and were segmented, respectively, into 2, 4 and 3

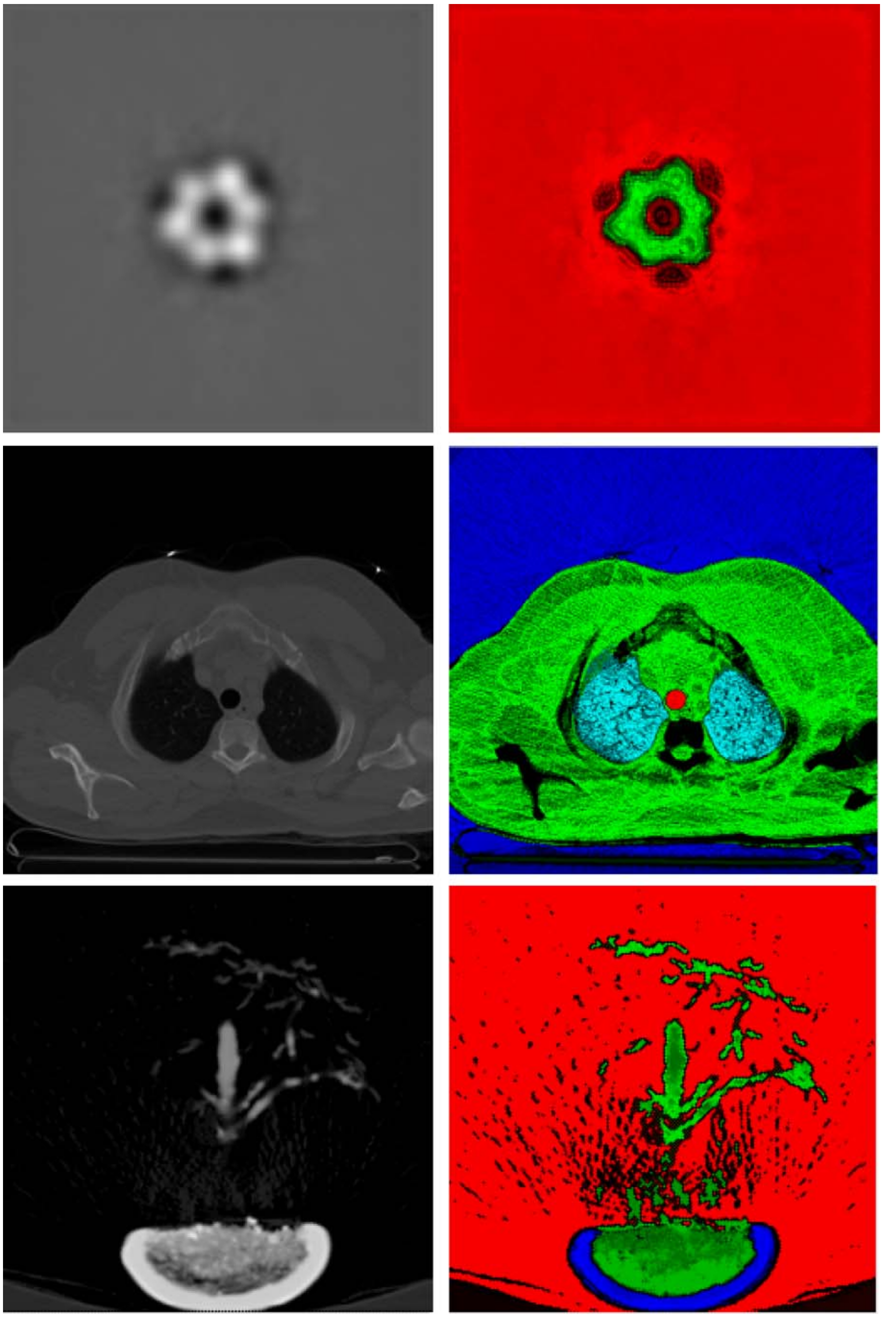

Fig. 4. Slices of an EM volume (top) and two CT volumes on the fcc grid are shown on the left. The corresponding slices of 2-, 4- and 3-segmentations of the respective volumes, obtained using the implementation of MOFS discussed in Section 5, are shown on the right. (All six images were interpolated for display purposes. This figure appears in color in the online version of the paper.)

objects. For these volumes, the speedup factors of the new implementation (as compared to the original implementation) were 6.71, 7.31 and 7.35, respectively.

6. Comparison with the approach of Udupa, Saha and Lotufo

In a recent paper [14] on the topic of segmentation of multiple objects using fuzzy connectedness, Udupa, Saha and Lotufo claimed that the theoretical results of [8], which have been restated in Section 2, are particular cases of the results described by them. We disagree with this claim, and in this section we present the reasons for our disagreement. These reasons fall into three categories:

1. our approach is more general than that of [14];
2. even in the special cases where both approaches are applicable, they behave differently: they produce different M-segmentations and our approach is inherently more efficient; and
3. the mathematical nature of our main result (Theorem 1) is quite different from anything presented in [14].

As opposed to our general approach, in [14] the only V's that are discussed are of the form (1) and M is restricted to be 2. The latter is justified on the basis that, for any one of the objects, all the other objects can be considered to be its "background" and so there is no loss of generality. We do not think that this justification is valid in all cases and, even when it is valid, it seems to us desirable to achieve simultaneous M-segmentations of the type illustrated in Figs. 2 and 3. In fact, in a recently published article [13], the first two authors of [14] state: "The iterative relative fuzzy connectedness theory and algorithms currently available are only for two objects. Considering the fact that most scenes contain more than two objects, for these algorithms to be useful, they and the associated theory should be extended to multiple objects."

To compare further our approach to that of [14] we need to make precise how objects are defined in [14]. Two different ways of defining objects are presented there.

The first way is called *relative fuzzy connectedness* (RFC). For its application it is assumed that $M = 2$, V is some set defined by (1), $\Psi = (\psi, \psi)$ for some reflexive and symmetric fuzzy spel affinity ψ such that V is ψ-connected, and both sets of seed spels V_1 and V_2 have exactly one element. Under these restrictions, RFC defines a 2-segmentation as follows. For $1 \leqslant m \leqslant 2$ and for any $c \in V$, let μ_m^c denote the ψ-strength of the strongest chain from (the unique element of) V_m to c. Then, let

$$\sigma_1^c = \begin{cases} \mu_1^c & \text{if } \mu_1^c > \mu_2^c, \\ 0 & \text{otherwise,} \end{cases} \tag{16}$$

$$\sigma_2^c = \begin{cases} \mu_2^c & \text{if } \mu_1^c \leqslant \mu_2^c, \\ 0 & \text{otherwise,} \end{cases} \tag{17}$$

and $\sigma_0^c = \max\{\sigma_1^c, \sigma_2^c\}$. Clearly, σ is a 2-semisegmentation of V. It is not difficult to prove, using the connectedness of V under the fuzzy spel affinity ψ, that σ is a 2-segmentation.

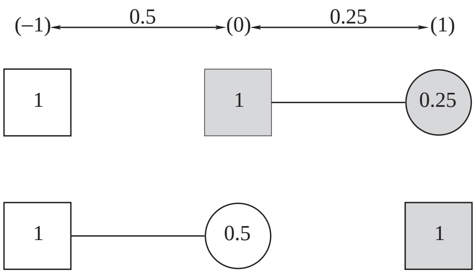

Fig. 5. Illustration of a desirable lack of “robustness” in the 2-segmentation determined by Theorem 1. The top row describes the 2-fuzzy graph by specifying the nonzero fuzzy spel affinities. In the other two rows there is shading for the object and no shading for the background, squares denote seed spels and circles denote other spels, the numbers indicate the grades of membership and the lines indicate the affinities which are responsible for the capture of the non-seed spel by the object (middle row) or by the background (bottom row). In the middle row (1) is captured by the object since the path from the background seed spel (-1) is blocked by the object seed spel (0). In the bottom row (0) is captured by the background, since its seed (-1) has a stronger fuzzy spel affinity to (0) than does the seed (1) of the object.

To illustrate this definition, consider the seeded 2-fuzzy graph $(\overline{V}, (\overline{\psi_1}, \overline{\psi_1}), (\overline{V_1}, \overline{V_2}))$. It is easy to see that the resulting 2-segmentation will be $\sigma^{(-1)} = (1, 0, 1)$, $\sigma^{(0)} = (1, 1, 0)$, and $\sigma^{(1)} = (0.25, 0, 0.25)$, the last due to the fact that $\mu_1^{(1)} = \mu_2^{(1)} = 0.25$. Note that this is different from the $\overline{\sigma}$ satisfying Theorem 1(i), for which $\overline{\sigma}^{(1)} = (0.25, 0.25, 0)$.

This illustrates that even if we restrict ourselves to that subset of connectable seeded M-fuzzy graphs to which RFC is applicable, there can be differences between the 2-segmentation produced by Theorem 1 and the one determined by RFC. We proceed to discussing this further.

For the example presented in the second paragraph above, the two 2-segmentations are essentially different: in $\overline{\sigma}$ (1) belongs to Object 1 and RFC tells us that it is in the background (Object 2). This is because RFC does not have the concept of “blocking” of the chain $\langle(-1), (0), (1)\rangle$ by the seed spel (0) of Object 1. We consider this to be a disadvantage of RFC (but this is more a matter of opinion than a supportable hypothesis).

RFC has a “robustness” property (Proposition 2.4 of [14]) which in our terminology can be restated as follows. If σ is the 2-segmentation defined by RFC and $\sigma_1^q > 0$, then if we replace V_1 by $\widetilde{V_1} = \{q\}$, we get by RFC a 2-segmentation $\widetilde{\sigma}$ such that, for all $c \in V$, $\sigma_1^c > 0$ if, and only if, $\widetilde{\sigma}_1^c > 0$. While it can indeed be argued that this is a desirable property (as it is done in [14]), there are situations where it seems to us to be counterproductive. For example, in the case considered above (in which $\overline{\sigma}_1^{(0)} = 1$ and $\overline{\sigma}_2^{(0)} = 0$) we had that $\overline{\sigma}_1^{(1)} = 0.25 > 0$. We find this quite acceptable. However, if we replace $\overline{V_1} = \{(0)\}$ by the set $\widetilde{V_1} = \{(1)\}$, then we get $\widetilde{\sigma}_1^{(0)} = 0$ and $\widetilde{\sigma}_2^{(0)} = 0.5$. This seems to us quite appropriate, even though it violates the robustness criterion of [14]; see Fig. 5.

Another difference is that the definition in RFC is not symmetric; if we interchange V_1 and V_2 that does not result in interchanging σ_1^c and σ_2^c (see the asymmetry in the definitions (16) and (17)). As a result of this, even though the “object” (Object 1 in our terminology) produced by RFC is guaranteed to be ψ-connected, the “background” (Object 2 in our terminology) is not. This is illustrated in the example $(\overline{V}, (\overline{\psi_1}, \overline{\psi_1}), (\overline{V_1}, \overline{V_2}))$ above, in which the background produced by RFC is disconnected. The *M*-segmentation defined by Theorem 1 is perfectly symmetric: if we permute the ψ_m's and the V_m's in the same way, then we will get exactly the corresponding permutation of the σ_m^c's (and the connectedness of all spels in an object to at least one seed spel of the object, as expressed by (13), will be preserved). We consider this also a disadvantage of RFC.

To overcome the lack of ability of RFC to achieve some desired results, [14] introduces a second method of object definition: *iterative relative fuzzy connectedness* (IRFC). Translated into our terminology, IRFC defines objects as follows.

Given a seeded 2-fuzzy graph $(V, (\psi, \psi), (V_1, V_2))$ (with all the previously stated restrictions in the approach of [14] implied), IRFC produces a sequence ${}^0\psi_2, {}^1\psi_2, \ldots$ of spel-adjacencies and a sequence of ${}^0\sigma, {}^1\sigma, \ldots$ of 2-segmentations defined as follows. ${}^0\psi_2 = \psi$ and ${}^0\sigma$ is the 2-segmentation defined by RFC. Now assume that, for some $i > 0$, we have already obtained ${}^{i-1}\psi_2$ and ${}^{i-1}\sigma$. For all $c, d \in V$, we define

$$
{}^i\psi_2(c, d) = \begin{cases} 1 & \text{if } c = d, \\ 0 & \text{if } {}^{i-1}\sigma_1^c > 0 \text{ or } {}^{i-1}\sigma_1^d > 0, \\ \psi(c, d) & \text{otherwise.} \end{cases} \tag{18}
$$

Using the notation ${}^i\psi_1 = \psi$, for all *i*, we define, for $1 \leqslant m \leqslant 2$ and for any $c \in V$, ${}^i\mu_m^c$ as the ${}^i\psi_m$-strength of the strongest chain (in *V*) from V_m to *c*. Then ${}^i\sigma$ is defined just as σ is defined in RFC using (16) and (17), but with μ_m^c replaced by ${}^i\mu_m^c$ everywhere. Note that the middle line of (18) causes the “blocking” of chains in Object 2 by spels in Object 1.

To illustrate this definition, consider again the seeded 2-fuzzy graph $(\overline{V}, (\overline{\psi_1}, \overline{\psi_1}), (\overline{V_1}, \overline{V_2}))$. Accordingly ${}^0\sigma^{(-1)} = (1, 0, 1)$, ${}^0\sigma^{(0)} = (1, 1, 0)$ and ${}^0\sigma^{(1)} = (0.25, 0, 0.25)$. Using (18), we see that ${}^1\psi_2(c, d) = 0$ if $c = (0)$ or $d = (0)$. This causes the value of ${}^1\mu_2^{(1)}$ to be 0 (while ${}^1\mu_1^{(1)} = 0.25$), and so ${}^1\sigma^{(1)} = (0.25, 0.25, 0)$; i.e., (1) gets assigned to Object 1 rather than to Object 2. Further iterations will not change the 2-segmentation; i.e., ${}^i\sigma = {}^1\sigma$,for $i \geqslant 1$. Note that in fact this ${}^1\sigma$ is the very $\overline{\sigma}$ that is determined by Theorem 1 for the same seeded 2-fuzzy graph.

However, this is not always the case; we now give an example in which the 2-segmentation determined by Theorem 1 is different from all the 2-segmentations produced by IRFC. Consider the seeded 2-fuzzy graph $(\overline{V}, (\psi, \psi), (\overline{V_1}, \overline{V_2}))$, where ψ is determined by $\psi((-1), (0)) = 1$, $\psi((0), (1)) = 0.5$, and $\psi((-1), (1)) = 0$. The 2-segmentation determined by Theorem 1 is $\sigma^{(-1)} = (1, 1, 1)$, $\sigma^{(0)} = (1, 1, 1)$ and $\sigma^{(1)} = (0.5, 0.5, 0.5)$. On the other hand, it is easy to see that, for all $i \geqslant 0$, the 2-segmentation provided by IRFC is ${}^i\sigma^{(-1)} = (1, 0, 1)$, ${}^i\sigma^{(0)} = (1, 0, 1)$ and ${}^i\sigma^{(1)} = (0.5, 0, 0.5)$. That is, Theorem 1 provides us with a 2-segmentation in which every spel belongs to both objects, but IRFC provides us with a 2-segmentation in which every spel is only in Object 2 (the background), including even the seed spel of Object 1. This seems to us a disadvantage of IRFC. This disadvantage becomes even more obvious if we do something that is allowed in [8] but not in the theory of [14], namely if we replace the seeded 2-fuzzy graph by $(\overline{V}, (\psi_1, \psi_2), (\overline{V_1}, \overline{V_2}))$where ψ_1 is the ψ defined above and ψ_2 is ψ except for $\psi_2((0), (1)) = 0.25$. Then, in our opinion quite reasonably, the 2-segmentation provided by Theorem 1 is $\sigma^{(-1)} = \sigma^{(0)} = (1, 1, 1)$ and $\sigma^{(1)} = (0.5, 0.5, 0)$; i.e., (1) belongs only to Object 1. The ability to achieve this depends on our freedom of selecting ψ_1 and ψ_2 independently and it cannot be imitated by IRFC as defined in [14].

An advantage of our approach over RFC is that we do *not* have to fully calculate the connectedness value of a spel c with respect to each of the objects to determine to which of those objects c belongs. RFC must do just that. In IRFC, an iterative algorithm, this disadvantage is compounded by the fact that the connectedness values for one of the objects are repeatedly recalculated until there are no further changes. In our approach, the connectedness values $\mu_{\sigma,m,V_m}(c)$ associated with all the objects are calculated simultaneously. Note that when we calculate a spel c's "potential connectedness value" with respect to an object that will later turn out not to contain c, we always discard the calculated value before it is ever used for calculating connectedness values of any other spels. In contrast, RFC uses the connectedness value of a spel with respect to each of the objects to calculate connectedness values of other spels with respect to that object. Thus our approach calculates connectedness values with less computational effort than RFC's approach does. Moreover, in the special (but frequently used) case in which V is of the form (1) and a fuzzy spel affinity is 0 unless c and d are adjacent, Step 15 of the MOFS algorithm and Step 14 of its fast implementation can be coded so that, most of the time, we do not waste computational effort in calculating connectedness values for spels for which σ_m^c will eventually be set to zero (as in (3) of our theorem and in (16)–(17) of the approach of [14]).

Finally, we comment on Theorem 1. Its general nature is the following. "Let G be a graph. A partial labeling of the nodes of G is said to have Property X if the label at each node can be determined from the labels assigned to the other nodes by a Procedure Y. We claim that for every graph G, there is one and only one partial labeling that has Property X, and it is in fact a total labeling provided G is connected." This is a result of some substance: there is no a priori reason to believe that, for all graphs, there would necessarily be a labeling with Property X, or that that labeling (if it exists) should be unique and/or total. It is trivially true that the deterministic algorithms RFC and IRFC produce unique labelings. It is much more difficult to prove that the labeling defined by the property of Theorem 1(i) exists and is unique.

7. Summary

This paper presents material associated with the first two authors' MOFS algorithm for simultaneous fuzzy segmentation of multiple objects that complements the material presented in [8]. In particular, it is proved that an M-semisegmentation satisfying the property stated in Theorem 1(i) exists, and it is clear from our proof of Theorem 1 that this desirable mathematical property uniquely characterizes the output of the MOFS algorithm. Additionally, our proof identifies a set of other properties (**A**, (13), and consistency of each pair of spels) that provide an alternative unique characterization of the output. A new, and usually considerably faster, implementation of MOFS is also presented. Our method of fuzzy segmentation is more general and efficient than the RFC and IRFC algorithms of [14].

References

[1] B.M. Carvalho, Cone-beam helical CT virtual endoscopy: reconstruction, segmentation and automatic navigation, Ph.D. Thesis, Computer and Information Science Program, University of Pennsylvania, 2003.

[2] B.M. Carvalho, E. Garduño, G.T. Herman, Multiseeded fuzzy segmentation on the face centered cubic grid, in: S. Singh, N. Murshed, W. Kropatsch (Eds.), Advances in Pattern Recognition: Second International Conference—ICAPR 2001 (LNCS 2013), Rio de Janeiro, Brazil, Springer, Berlin, 2001, pp. 339–348.

[3] B.M. Carvalho, C.J. Gau, G.T. Herman, T.Y. Kong, Algorithms for fuzzy segmentation, Pattern Anal. Appl. 2 (1999) 73–81.

[4] B.M. Carvalho, G.T. Herman, T.Y. Kong, Simultaneous fuzzy segmentation of multiple objects, in: A. Del Lungo, V. Di Gesù, A. Kuba (Eds.), Electronic Notes in Discrete Mathematics, vol. 12, Elsevier Science Publishers, 2003.

[5] T.H. Cormen, C.E. Leiserson, R.L. Rivest, Introduction to Algorithms, MIT Press, Cambridge, MA, 1990.

[6] R.O. Duda, P.E. Hart, Pattern Classification and Scene Analysis, Wiley, New York, 1973.

[7] G.T. Herman, Geometry of Digital Spaces, Birkhäuser, Boston, MA, 1998.

[8] G.T. Herman, B.M. Carvalho, Multiseeded segmentation using fuzzy connectedness, IEEE Trans. on Pattern Anal. Mach. Intell. 23 (2001) 460–474.

[9] A.K. Jain, M.N. Murty, P.J. Flynn, Data clustering: a review, ACM Comput. Surveys 31 (1999) 264–323.

[10] H.A. Moghaddam, J.F. Lerallut, Volume visualization of the heart using MRI 4D cardiac images, J. Comput. Inform. Tech. 6 (1998) 215–228.

[11] L.G. Nyul, A.X. Falcão, J.K. Udupa, Fuzzy-connected 3D image segmentation at interactive speeds, Graph. Models 64 (2002) 259–281.

[12] A. Rosenfeld, Fuzzy digital topology, Inform. and Control 40 (1979) 76–87.

[13] J.K. Udupa, P.K. Saha, Fuzzy connectedness and image segmentation, Proc. IEEE 91 (2003) 1649–1669.

[14] J.K. Udupa, P.K. Saha, R.A. Lotufo, Relative fuzzy connectedness and object definition: theory, algorithms, and applications in image segmentation, IEEE Trans. on Pattern Anal. Mach. Intell. 24 (2002) 1485–1500.

[15] J.K. Udupa, S. Samarasekera, Fuzzy connectedness and object definition: theory, algorithms and applications in image segmentation, Graph. Models Image Proc. 58 (1996) 246–261.

[16] C.T. Zahn, Graph-theoretic methods for detecting and describing Gestalt clusters, IEEE Trans. Comp. 1 (1971) 68–86.