Advances in Constrained Connectivity

Pierre Soille and Jacopo Grazzini

Spatial Data Infrastructures Unit
Institute for Environment and Sustainability
Joint Research Centre of the European Commission
T.P. 262, via Fermi 2749, I-21027 Ispra, Italy
{Pierre.Soille,Jacopo.Grazzini}@jrc.it

Abstract. The concept of constrained connectivity [Soille 2008, PAMI] is summarised. We then introduce a variety of measurements for characterising connected components generated by constrained connectivity relations. We also propose a weighted mean for estimating a representative value of each connected component. Finally, we define the notion of spurious connected components and investigate a variety of methods for suppressing them.

1 Introduction

A segmentation of the definition domain X of an image is usually defined as a partition of X into disjoint connected subsets $X_i, \ldots, X_n$ (called segments) such that there exists a logical predicate P returning true on each segment but false on any union of adjacent segments [1]. That is, a series of subsets X_i of the definition domain X of an image forms a segmentation of this image if and only if the following four conditions are met (i) $\cup_i(X_i) = X$, (ii) $X_i \cap X_j = \emptyset$ for all $i \neq j$, (iii) $P(X_i) =$ true for all i, and (iv) $P(X_i \cup X_j) =$ false if X_i and X_j are adjacent.

With this classical definition of image segmentation, given an arbitrary logical predicate, there may exist more than one valid segmentation. For example, the logical predicate returning true on segments containing one and only one regional minimum and false otherwise lead to many possible segmentations. The watershed transformation definition considers the additional constraint that there should exist a steepest slope path linking each pixel of the segment to its corresponding minimum for the logical predicate to return true. Still, this does not guarantee that there is a unique solution because the steepest slope path of a pixel is not necessarily unique (problem of ties).

If uniqueness of the result is required, logical predicates based on equivalence relations should be considered[1]. Indeed, it has been known for a long time that there exists a one-to-one correspondence between the partitions of a set and the equivalence relations on it, e.g., [2, p. 48]. Since connectivity relations are equivalence relations, logical predicates based on connectivity relations naturally lead

[1] A binary relation which is reflexive, symmetric, and transitive is called an equivalence relation.

D. Coeurjolly et al. (Eds.): DGCI 2008, LNCS 4992, pp. 423–433, 2008.

to unique segmentations. For example, the trivial connectivity relation stating that two pixels are connected if and only if they can be joined by an iso-intensity path breaks digital images into segments of uniform grey scale [3]. They are called plateaus in fuzzy digital topology [4] and flat zones in mathematical morphology [5]. In most cases, the equality of grey scale is a too strong homogeneity criterion so that it produces too many segments. Consequently, the resulting partition is *too fine*. A weaker connectivity relation consists in stating that two pixels of a grey tone image are connected if there exists a path of pixels linking these pixels and such that the grey level difference along adjacent pixels of the path (i.e., weights of the edges of the path) does not exceed a given threshold value. In this paper, we call this threshold value the local range parameter and denote it by α. Accordingly, we call the resulting connected components the α-connected components. This idea was introduced in image processing by Nagao *et al.* in the late seventies [6]. The resulting connected components are called quasi-flat zones [7] in mathematical morphology. The concept of α-connected components predates developments in image processing since it is at the very basis of the single linkage clustering method [8]. Although α-connected components often produce adequate image partitions they fail to do so when distinct image objects (with variations of intensity between adjacent pixels not exceeding α) are separated by one or more transitions going in steps having an intensity height less than or equal to α. Indeed, in this case, these objects appear in the same α-connected component so that the resulting partition is *too coarse*.

A natural solution to this problem is to limit the difference between the maximum and minimum values of each connected component by introducing a second threshold value called hereafter global range parameter and denoted by ω. This idea has been originally introduced in [9]. However, the relation at the basis of the developments of [9] is not an equivalence relation because it is not transitive and therefore does not guarantee the generation of unique connected components. This problem has been solved in [10] by introducing the notion of constrained connectivity. In the present paper, we expand on the results of [10].

The paper is organised as follows. Constrained connectivity relations originally proposed in [10] are summarised in section 2. We then propose in Sec. 3 a series of measurements that can be applied to each connected component and introduce the notion of local connectivity index leading to a weighted mean of its intensity values. Image segmentation based on constrained connectivity is studied in section 4. Particular emphasis is given on the analysis of small and usually undesirable segments. We suggest several procedures for suppressing them while preserving the hierarchical properties of partitions based on constrained connectivity. Experiments conducted on a benchmark aerial image are discussed in section 5.

2 Constrained Connectivity Relations

After a reminder about the well established notion of alpha-connectivity, this section summarises the notion of constrained connectivity recently introduced in [10].

2.1 Alpha-Connectivity

Two pixels p and q of an image f are α-connected if there exists a path going from p to q such that the range of the intensity values between two successive pixels of the path does not exceed the value of the local range parameter α [6]. By definition, a pixel is α-connected to itself. Accordingly, the α-connected component of a pixel p is defined as the set of image pixels that are α-connected to this pixel. We denote this connected component by α-CC(p):

$$\alpha\text{-CC}(p) = \Big\{p\Big\} \cup \Big\{q \;\Big|\; \text{there exists a path } \mathcal{P} = (p = p_1, \ldots, p_n = q),\ n > 1, \\ \text{such that } \mathsf{R}\{f(p_i), f(p_{i+1})\} \leq \alpha \text{ for all } 1 \leq i < n\Big\},$$

where the range function R calculates the difference between the maximum and the minimum values of a nonempty set of intensity values.

More restrictive connectivity relations detailed in Sec. 2.2 exploit the total ordering relation between the α-connected components of a pixel. Indeed, for all local range parameters α less than or equal to a given local range parameter α', the α-connected component of a pixel p is included in the α'-connected component of this pixel:

$$\alpha\text{-CC}(p) \subseteq \alpha'\text{-CC}(p) \text{ for all } \alpha \leq \alpha'. \tag{1}$$

This hierarchy is known since the fifties in the field of combinatorial optimisation, see [11] for a detailed survey till 1960. Indeed, it is at the root of the greedy algorithm of Kruskal [12] for solving the minimum spanning tree problem. In this algorithm, referred to as 'construction A' in [12], the edges of the graph are initially sorted by increasing edge weights. Then, the minimum spanning tree T is defined recursively as follows: the next edge is added to T if and only if together with T it does not form a circuit. That is, assuming the edge weights are defined by the range of the intensity values of the two nodes (pixels) they link, there is a one-to-one correspondence between (i) the α-connected components and (ii) the subtrees obtained for a distance α in Kruskal's greedy solution to the minimum spanning tree problem. This hierarchy of subtrees is itself at the very basis of the dendrogram representation of the single linkage clustering [8]. This clustering method was put forward by Sneath [13] as a convenient way of summarising taxonomic relationships in the form of taxonomic trees also called similarity trees or dendrograms.

2.2 (Alpha,Omega)-Connectivity

We define the (α, ω)-connected component of an arbitrary pixel p as the largest α_i-connected component of p such that (i) $\alpha_i \leq \alpha$ and (ii) its range is lower than or equal to ω [10]:

$$(\alpha, \omega)\text{-CC}(p) = \bigvee \Big\{\alpha_i\text{-CC}(p) \;\Big|\; \alpha_i \leq \alpha \text{ and } \mathsf{R}\Big(\alpha_i\text{-CC}(p)\Big) \leq \omega\Big\}. \tag{2}$$

The existence of a largest α_i-connected component is secured thanks to the total order relation between the α_i-connected components of a pixel (Eq. 1). Two pixels p and q are (α, ω)-connected if and only if $q \in (\alpha, \omega)$-CC(p).

Beyond range parameters, one may consider other constraints such as a connectivity index indicating the degree of cohesion of each α-connected component. This idea leads to the notion of α-strong connectivity detailed in [10]. A further generalisation to arbitrary logical predicates is presented in [14].

3 Connected Component Representation

First, a series of useful measurements that can be applied to each connected component is presented. We then indicate a method for computing a representative value for each connected component, taking into account their internal cohesion.

3.1 Measurements

Measurements performed on each connected component provide us with a set of features useful for classification purposes and subsequent processing. We propose the definition of the difference image Δ_A mapping the difference between α and the maximum value of α_i leading to the (α, ω)-connected component of p:

$$\Delta_A[\mathrm{CC}(p)] = \alpha - \max\left\{\alpha_i \mid \alpha_i\text{-CC}(p) = (\alpha, \omega)\text{-CC}(p)\right\}.$$

Similarly, the difference image Δ_Ω measures the difference between ω and the actual range of the connected component:

$$\Delta_\Omega[\mathrm{CC}(p)] = \omega - \mathsf{R}\Big((\alpha, \omega)\text{-CC}(p)\Big).$$

The difference between the maximum and minimum value of α_i leading to the (α, ω)-connected component is proportional to the strength of the external isolation of the component:

$$\max\left\{\alpha_i \mid \alpha_i\text{-CC}(p) = (\alpha, \omega)\text{-CC}(p)\right\} - \min\left\{\alpha_i \mid \alpha_i\text{-CC}(p) = (\alpha, \omega)\text{-CC}(p)\right\}.$$

The connectivity index function [10] obtained for increasing threshold range values could also be used as feature vector characterising each connected component. It is used in the following section for calculating a representative value of each connected component.

3.2 Representative Value

Within the scope of image simplification, one needs to estimate a representative value for each connected segment[2]. This is also necessary when iterating the

[2] For the estimation of a representative value within geodesic adaptive neighbourhood instead of connected components, see [15] in this volume.

partitioning procedure. In this latter case, the estimated values at each step of the iteration influence the segments obtained in the successive steps.

A common choice for the representative value of a segment is the average of the grey levels of the pixels belonging to it. Some approaches also propose to select it as the local mode of the segment [16].

We propose here to associate with each segment a weighted mean of the intensity values of the pixels of the segment. This is achieved by calculating for each pixel of the segment the number of adjacent pixels that belong to this segment and that are within a range of α. More generally, rather than looking for adjacent pixels only, one can analyse larger neighbourhoods (of size n) to estimate a representative value $\psi_n[\mathrm{CC}(p)]$ of any segment $\mathrm{CC}(p)$:

$$\psi_n[\mathrm{CC}(p)] = \left\{ \frac{\sum \mathsf{CI}_n(p_i) f(p_i)}{\sum \mathsf{CI}_n(p_i)} \;\middle|\; p_i \in \alpha\text{-}\mathrm{CC}(p) \right\}, \tag{3}$$

with the *local connectivity index* of order n of the pixel p defined as $\mathsf{CI}_n(p) = \mathrm{card}(p_i \in \mathrm{CC}(p))$ such that there exists a path $\mathcal{P}$ with length n and $\mathsf{R}(f(p), f(p_i)) \leq \alpha$. This idea is related to the concept of (global) connectivity index function as introduced in [10] and defined at the level of a connected component.

With this definition, we have in particular $\mathsf{CI}_1[p] \leq N$ where N is imposed by the graph connectivity definition (4 or 8 in the square grid). This way, in the computation of the weighted mean ψ_1, large weights are assigned to the 'core' pixels of the segment, with large connectivity index, while lower weights are assigned to pixels with smaller connectivity index. Notice moreover that the border pixels of the segment, that have at least one connection with a pixel from another segment, have a connectivity index automatically forced to a value $\mathsf{CI}_n[p] < N$. Thus, pixels lying on the internal segment boundaries are assigned lower weights in the weighting procedure defined by Eq. 3 and they contribute less to the final grey level of the segment ($\psi_1[\mathrm{CC}(p)]$).

4 Partition Filtering

Constrained connectivity relations partitions the image definition domain into labelled connected components. In addition, by varying the threshold values of the constraints, partition of increasing coarseness degree are obtained. This idea is applied to hierarchical image decomposition and simplification in [10]. Interestingly, any level of the hierarchy can be directly computed without requiring knowledge of the previous levels, in contrast to most alternative partition hierarchies [17].

The generated partitions deliver puzzle pieces that can be further assembled depending on application dependent rules. However, by essence, the method does not take any size criterion into account. It follows that regions as small as one pixel may survive even for large values of the constraint threshold values. We study hereafter the origin of these small regions and propose some approaches for suppressing them in cases this is required by the application at hand.

4.1 Characterisation of Small Regions

We define small regions as regions that cannot contain the elementary structuring element defined by a pixel and its adjacent neighbours (4- or 8-neighbours in the square grid). They are extracted by the following 3-step procedure:

1. perform the union of the erosion of each connected component of the labelled partition. This can be achieved by initialising the output image to 1 and then scan the input image while checking for each position that the structuring element centred at this position covers pixels with the same label value. If this is not the case, the value of the output image at the current position is set to 0. The resulting binary image is then multiplied by the input labelled partition. This image corresponds to the union of the erosion of each connected component of the labelled partition and is referred to as the *marker image* hereafter;
2. reconstruct the labelled partition from the marker image (reconstruction operation on labels, that is, the markers propagate only within the region having the same label value as the marker);
3. define *small regions* as the arithmetic difference between the initial labelled partition and the reconstructed partition as per step 2.

The resulting small segments are then categorised into two classes having different origins:

– 1 pixel thick segments containing at least one regional extremum (union of regional minima and maxima). These regions may either be due to noise or thin relevant structures.
– 1 pixel thick segments that do not contain any regional extremum. These regions are usually due to the limited resolution of the digital image leading to non ideal step edges spanning over 1 or more pixels. We call them *transition regions* (for alternative definitions of transition regions, see [18,19]). Often, transition regions are located at the boundaries between larger regions.

In both situations, these small regions cannot grow further because their growth would lead to a violation of either the local or global range constraints. We explore hereafter a number of methods for reducing or even suppressing small regions of one or both types.

4.2 Filtering Procedures

Filters can be applied either before or after the computation of the constrained connected components.

Pre-filtering. Filters reducing irrelevant local variations can help aggregating small regions into larger regions since the largest α_i-connected component satisfying the input constraints will usually be obtained for larger values of α_i. The occurrence of the first type of small regions can reduced by preprocessing the image with a filter removing isolated pixels. If necessary, more active filters such as the self-dual reconstruction of the input image from its median filter or self-dual area filters [20] can be considered.

Post-filtering. The filtering procedure consists in computing a partition given local and global range parameters. Transition regions defined as small regions not containing any regional extrema are then extracted and considered as *spurious regions*. The resulting gaps are then filled using a seeded region growing algorithm [21] (see also [22] for a version suitable for connected operators and multispectral images). This procedure ensures that all unwanted regions are suppressed but at the cost of some arbitrary decisions unless the ties are tracked (see seeded region growing algorithm enhanced in [23] to address order dependence issues).

5 Experiments

Figure 1a shows an aerial image retrieved from the miscellaneous section of image database of the University of Southern California (USC-SIPI Image Database). Figure 1b shows the partition obtained using the same value (64 grey levels) for

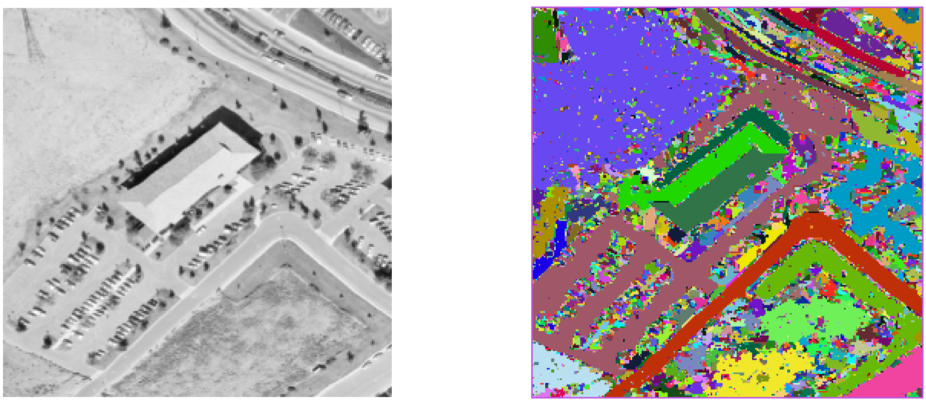

(a) 256 × 256 aerial image (54,364 iso-intensity connected components).

(b) (α, ω)-partition with $\alpha = \omega = 64$ (8,664 regions).

Fig. 1. Input aerial image and resulting partition using local and global range thresholds equal to 64. The input image corresponds to the lower left quarter of the image 5.2.09 of the miscellaneous section of the USC-SIPI image database, see http://sipi.usc.edu/database/

the local and global range parameters and considering the 4-connected graph. A simplified image can be generated by setting each region to the weighted mean as proposed in Sec. 3.2 leads to the simplified image shown in Fig. 2a. A comparison with the image obtained using the non weighted mean is shown in Fig. 2b. This image reveals that the weighted mean generate a more contrasted image even if differences can be hardly perceived through a visual comparison between mean and weighted mean representations.

Figure 3 illustrates the post-filtering procedure whereby transition regions are suppressed (1 pixel thick regions not containing any regional extremum). The

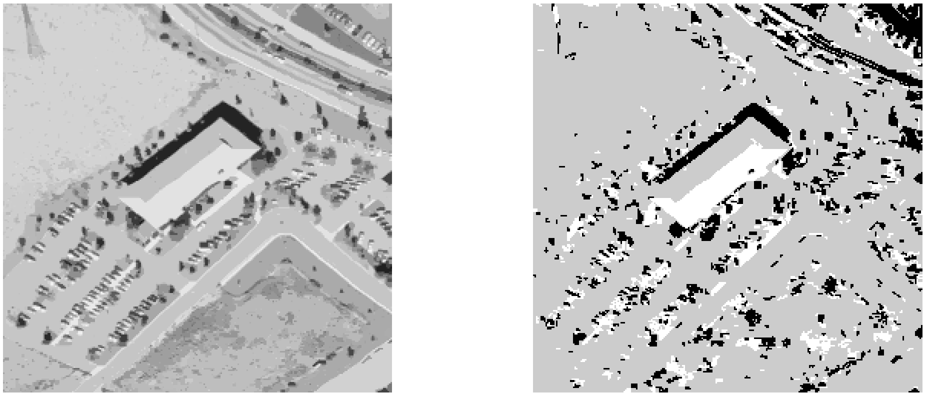

Fig. 2. Left: Edge preserving simplification of the image shown in Fig. 1a using the partition shown in Fig. 1b and weighted mean based on local connectivity index with adjacent pixels only (equation 3 with $n = 1$). Right: comparison between mean and weighted mean with grey for identical values, black for lower values with weighted mean, and white for greater values with weighted mean. In both cases, the mean was rounded to its integer part.

partition without transitions regions contains 2,897 regions contrary to the 8,664 regions of the initial partition (compare Fig. 3f with Fig. 1b by zooming on the electronic version).

6 Concluding Remarks and Perspectives

The concept of constrained connectivity offers a fruitful framework for creating image partitions and edge preserving filtering (image simplification). A non-exhaustive list of measurements characterising the generated connected components has been proposed. The use of these measurements for classification purposes will be reported in a follow-up paper together with their extension to multichannel images since the concept of constrained connectivity can be extended to these images [10]. The proposed notion of local connectivity index allows for the definition of a weighted mean for estimating a representative value of each connected component. Further improvements regard the estimation of a representative value for each connected component. For instance, rather than using the input local range threshold value α in the definition of the local connectivity index CI_n, the actual local range threshold value α_i (Eq. 2) that varies from one connected component to another could be considered.

We have also analysed the origin of the small connected components and categorised them into two main categories. A technique for removing spurious small regions corresponding to transition regions has been proposed and allows for a drastic reduction in the number of regions of the segmented images. Other techniques based on iterative methods could also be easily designed. Finally,

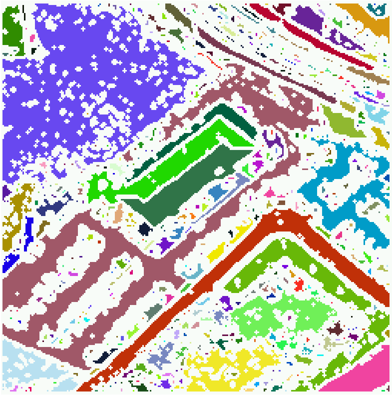

(a) Union of the erosion of CC of Fig. 1b.

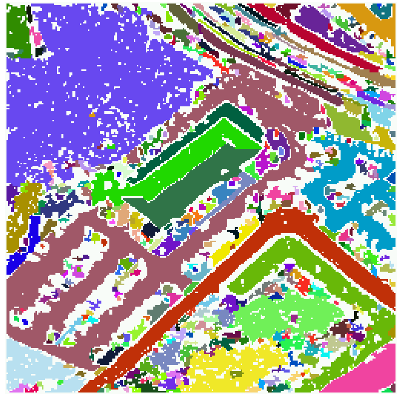

(b) Reconstruction of Fig. 1b from *(a)*.

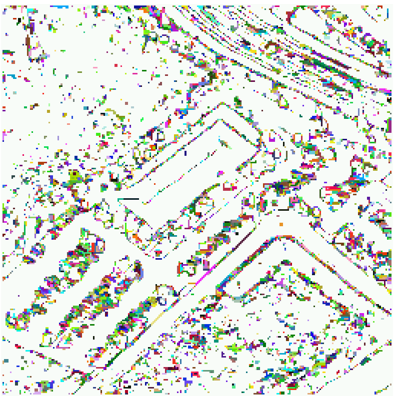

(c) Resulting small regions.

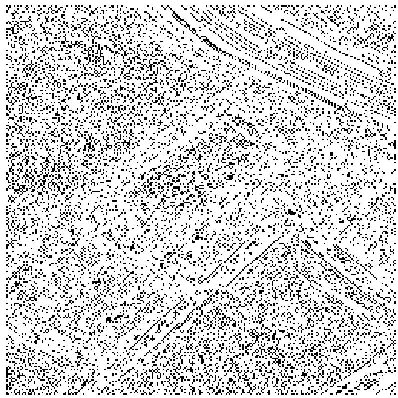

(d) Regional extrema of Fig. 1a.

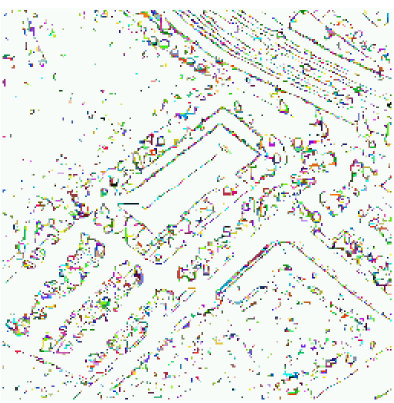

(e) Transition regions (small regions not containing a regional extremum).

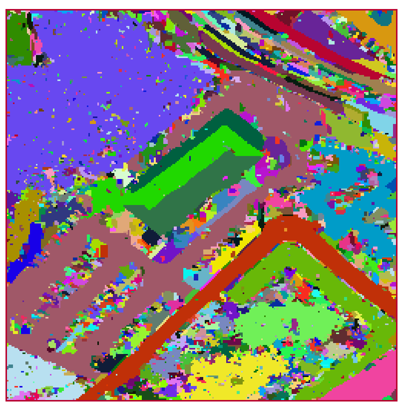

(f) Final partition with transition regions removed (2,897 regions).

Fig. 3. Partition filtering based on the removal of transition regions (see details in text)

comparisons with related hierarchical segmentation techniques [24,17] will be addressed in an extended version of this paper.

References

1. Zucker, S.: Region growing: childhood and adolescence. Computer Graphics and Image Processing 5, 382–399 (1976)
2. Jardine, N., Sibson, R.: Mathematical Taxonomy. Wiley, London (1971)
3. Brice, C., Fennema, C.: Scene analysis using regions. Artificial Intelligence 1, 205–226 (1970)
4. Rosenfeld, A.: Fuzzy digital topology. Information and Control 40, 76–87 (1979)
5. Serra, J., Salembier, P.: Connected operators and pyramids. In: Dougherty, E., Gader, P., Serra, J. (eds.) Image Algebra and Morphological Image Processing IV, vol. SPIE-2030, pp. 65–76 (1993)
6. Nagao, M., Matsuyama, T., Ikeda, Y.: Region extraction and shape analysis in aerial photographs. Computer Graphics and Image Processing 10, 195–223 (1979)
7. Meyer, F., Maragos, P.: Morphological scale-space representation with levelings. In: Nielsen, M., Johansen, P., Fogh Olsen, O., Weickert, J. (eds.) Scale-Space 1999. LNCS, vol. 1682, pp. 187–198. Springer, Heidelberg (1999)
8. Gower, J., Ross, G.: Minimum spanning trees and single linkage cluster analysis. Applied Statistics 18, 54–64 (1969)
9. Hambrusch, S., He, X., Miller, R.: Parallel algorithms for gray-scale image component labeling on a mesh-connected computer. In: Proc. of the 4th ACM Symposium on Parallel Algorithms and Architectures, pp. 100–108 (1992)
10. Soille, P.: Constrained connectivity for hierarchical image partitioning and simplification. IEEE Transactions on Pattern Analysis and Machine Intelligence (2008), [Available as online preprint since 18 October 2007]
11. Schrijver, A.: On the history of combinatorial optimization (till 1960). In: Aardal, K., Nemhauser, G., Weismantel, R. (eds.) Handbook of Discrete Optimization, pp. 1–68. Elsevier, Amsterdam (2005)
12. Kruskal, J.: On the shortest spanning subtree of a graph and the traveling salesman problem. Proceedings of the American Mathematical Society 7, 48–50 (1956)
13. Sneath, P.: Computers in taxonomy. Journal of General Microbiology 17, 201–226 (1957)
14. Soille, P.: On genuine connectivity relations based on logical predicates. In: Proc. of 14th Int. Conf. on Image Analysis and Processing, Modena, Italy, pp. 487–492. IEEE Computer Society Press, Los Alamitos (2007)
15. Grazzini, J., Soille, P.: Adaptive morphological filtering using similarities based on geodesic time. In: Coeurjolly, D., et al. (eds.) DGCI 2008. LNCS, vol. 4992, pp. 519–528. Springer, Heidelberg (2008)
16. Comaniciu, D., Meer, P.: Mean shift: a robust approach toward feature space analysis. IEEE Transactions on Pattern Analysis and Machine Intelligence 24, 603–619 (2002)
17. Kropatsch, W., Haxhimusa, Y.: Grouping and segmentation in a hierarchy of graphs. In: Bouman, C., Miller, E. (eds.) Proc. of the 16th IS&T SPIE Annual Symposium, Computational Imaging II, vol. SPIE-5299, pp. 193–204 (2004)
18. Tanimoto, S., Pavlidis, T.: The editing of picture segmentations using local analysis of graphs. Communications of the ACM 20, 223–229 (1977)

19. Crespo, J., Schafer, R., Serra, J., Gratin, C., Meyer, F.: The flat zone approach: a general low-level region merging segmentation method. Signal Processing 62, 37–60 (1997)
20. Soille, P.: Beyond self-duality in morphological image analysis. Image and Vision Computing 23, 249–257 (2005)
21. Adams, R., Bischof, L.: Seeded region growing. IEEE Transactions on Pattern Analysis and Machine Intelligence 16, 641–647 (1994)
22. Brunner, D., Soille, P.: Iterative area filtering of multichannel images. Image and Vision Computing 25, 1352–1364 (2007)
23. Mehnert, A., Jackway, P.: An improved seeded region growing. Pattern Recognition Letters 18, 1065–1071 (1997)
24. Guigues, L., Le Men, H., Cocquerez, J.P.: The hierarchy of the cocoons of a graph and its application to image segmentation. Pattern Recognition Letters 24, 1059–1066 (2003)