

Statistical Region Merging

Richard Nock and Frank Nielsen

Abstract—This paper explores a statistical basis for a process often described in computer vision: image segmentation by region merging following a particular order in the choice of regions. We exhibit a particular blend of algorithmics and statistics whose segmentation error is, as we show, limited from both the qualitative and quantitative standpoints. This approach can be efficiently approximated in linear time/space, leading to a fast segmentation algorithm tailored to processing images described using most common numerical pixel attribute spaces. The conceptual simplicity of the approach makes it simple to modify and cope with hard noise corruption, handle occlusion, authorize the control of the segmentation scale, and process unconventional data such as spherical images. Experiments on gray-level and color images, obtained with a short readily available C-code, display the quality of the segmentations obtained.

Index Terms—Grouping, image segmentation.

1 INTRODUCTION

IT is established since the Gestalt movement in psychology that perceptual grouping plays a fundamental role in human perception. Even though this observation is rooted in the early part of the 20th century, the adaptation and automation of the segmentation (and, more generally, grouping) task with computers has remained so far a tantalizing and central problem for image processing. Vision is widely accepted as an inference problem, i.e., the search of what caused the observed data [1]. In this respect, the grouping problem can be roughly presented as the transformation of the collection of pixels of an image into a visually meaningful partition of regions and objects.

This postulates implicitly the existence of optimal segmentation(s) which we should aim at recovering or approximating, and this task implies casting the perceptual formulation of optimality into a formalized, well-defined problem. A prominent trend in grouping focuses on graph cuts, mapping image pixels onto graph vertices, and the spatial relationships between pixels onto weighted graph edges. The objective is to minimize a cut criterion, given that any cut on this graph yields a partition of the image into (hopefully) coherent visual patterns. Cut criteria range from conventional [2] to more sophisticated criteria, tailored to grouping [3], [4], [5]. These are basically global criteria; however, the strategies adopted for their minimization range through a broad spectrum, from local [6] to global optimization [5], through intermediate choices [7], [3]. Global optimization strategies have the advantage to directly tackle the problem as a whole, and may offer good approximations [5], at possible algorithmic expenses though [3], [5].

In this paper, we focus on a different strategy which belongs to the family of region growing and merging techniques [8], [9]. In region merging, regions are sets of

pixels with homogeneous properties and they are iteratively grown by combining smaller regions or pixels, pixels being elementary regions. Region growing/merging techniques usually work with a statistical test to decide the merging of regions [9]. A merging predicate uses this test, and builds the segmentation on the basis of (essentially) local decisions. This locality in decisions has to preserve global properties, such as those responsible for the perceptual units of the image [8]. In Fig. 1, the grassy region below the castle is one such unit, even when its variability is high compared to the other regions of the image. In that case, a good region merging algorithm has to find a good balance between preserving this unit and the risk of overmerging for the remaining regions. Fig. 1b shows the result of our approach. As long as the approach is greedy, two essential components participate in defining a region merging algorithm: the merging predicate and the order followed to test the merging of regions. There is a lack of theoretical results on the way these two components interact together, and can benefit from each other. This might be partially due to the fact that most approaches use assumptions on distributions, more or less restrictive, which would make any theoretical insight into how region merging works restricted to such settings and, therefore, of possibly moderate interest (see, e.g., [10] for related criticisms).

Our aim in this paper is to propose a path and its milestones from a novel model of image generation, the theoretical properties of possible segmentation approaches to a practical, readily available system of image segmentation, and its extensions to miscellaneous problems related to image segmentation. First, the key idea of this model is to really formulate image segmentation as an inference problem [1]. It is the reconstruction of regions on the observed image, based on an unknown theoretical (true) image on which the true regions we seek are statistical regions whose borders are defined from a simple axiom. Second, we show the existence of a particular blend of statistics and algorithmics to process observed images generated with this model, by region merging, with two statistical properties. With high probability, the algorithm suffers only one source of error for image segmentation: overmerging, that is, the fact that some observed region may contain more than one true region. The algorithm does not suffer neither undermerging, nor the—most frequent—hybrid cases where observed regions may partially span several true regions. Yet, there is more: With

• R. Nock is with the Université Antilles-Guyane, Département Scientifique Inter-facultaire/GRIMAAG Lab., B.P. 7209, 97278 Schoelcher, Martinique, France. E-mail: rnock@martinique.univ-ag.fr.

• F. Nielsen is with Sony Computer Science Laboratories Inc., 3-14-13 Higashi Gotanda, Shinagawa-Ku, Tokyo 141-0022, Japan. E-mail: Frank.Nielsen@acm.org.

Manuscript received 8 Aug. 2003; revised 26 Jan. 2004; accepted 1 Apr. 2004. Recommended for acceptance by R. Basri.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0219-0803.



Fig. 1. An **RGB** image and the segmentation found by our segmentation method (regions are white bordered and averaged inside).

high probability, this overmerging error is, as we show, formally small as the algorithm manages an accuracy in segmentation close to the optimum, up to low order terms. The algorithm has some desirable features: It relies on a simple interaction between a merging predicate easily implementable, and an order in merging approximable in linear time. Furthermore, it can be adapted to most numerical feature description spaces (**RGB**, *HSI*, $L^*u^*v^*$, etc.). Third, we provide a C-code implementation of this last algorithm, which is a few hundred lines of C, and experiments on various benchmark images, as well as comparisons with other algorithms. Last, we show how to extend the algorithm to naturally cope with hard noise and/or significantly occluded images at very affordable algorithmic complexity. Though running the algorithm does not require tuning its parameters, the control of a statistical complexity parameter makes it possible to adjust the segmentation scale in a simple manner.

The next section presents our model of image generation. Section 3 details our analysis and algorithm, first for the gray-level setting, and then for color images. Section 4 presents our experiments. The last two sections conclude and detail the code availability.

2 PRELIMINARY NOTATIONS AND MODELS

The notation $|\cdot|$ stands for cardinal. The observed image, I , contains $|I|$ pixels, each containing **Red-Green-Blue** (**RGB**) values, each of the three belonging to the set $\{1, 2, \dots, g\}$ (in practice, we would have $g = 256$). We have deliberately




Fig. 3. Schematic view of some theoretical image I^* , and a corresponding observed image I . Only the average over **R, G, B** of the theoretical pixels' first moments are shown in I^* . According to the homogeneity property (see text), (a) shows the true (optimal) segmentation of I . a is the process generating the observed image (see also Fig. 2), and b is grouping's objective (i.e., find the statistical regions' borders, given I).

chosen not to use complex formulations of the colors, such as the $L^*u^*v^*$ space [10].

I is an observation of a perfect scene I^* we do not know of, in which pixels are perfectly represented by a family of *distributions*, from which each of the observed color channel is sampled. In I^* , the optimal (or true, or *statistical*) regions represent theoretical objects sharing a common *homogeneity property*:

- Inside any statistical region and given any color channel $\in \{\mathbf{R}, \mathbf{G}, \mathbf{B}\}$, the statistical pixels have the same expectation for this color channel.
- The expectations of adjacent statistical regions are different for at least one color channel $\in \{\mathbf{R}, \mathbf{G}, \mathbf{B}\}$.

I is obtained from I^* by sampling each statistical pixel for observed **RGB** values. Fig. 2 presents an example of a color channel for one pixel in I^* and how to generate the corresponding observed color channel of the pixel in I . In each pixel of I^* , each color channel is replaced by a set of exactly Q independent random variables (r.v.), taking positive values on domains bounded by g/Q , such that any possible sum of outcomes of these Q r.v. belongs to $\{1, 2, \dots, g\}$. Fig. 3 gives an example of some true image I^* (in fact, it is the result of our algorithm, ran on Fig. 3b) which displays the expectation of statistical pixels, and the observed image I generated from I^* . Given the homogeneity property, frontiers between true regions are connecting pixels with differences in their color expectations, and the ideal segmentation of I relies on the frontiers between the statistical regions shown on I^* in Fig. 3.

The sampling of each pixel and its color channels are supposed independent from each other. It is important to note that this is the only assumption we make on I^* , and the frequent independent identically distributed (i.i.d.) assump-




Fig. 2. Generation of a single color channel for one pixel from a statistical region O of I^* to some observed pixel of I .

tion is relaxed in this model to that of ordinary independence. Inside a statistical region, it can be the case that all distributions associated to each pixel are different, as long as the homogeneity property is satisfied. This freedom has a counterpart, which led us to introduce Q , not necessarily to make our model more general, but, essentially, for practical purposes: The conventional choice $Q = 1$ would actually make it hard to estimate reliably anything for small regions or, equivalently, would make it necessary to consider very large images to improve the statistical accuracy of the segmentation. Notice that Q is a parameter which makes sense: It allows us to quantify the *statistical complexity* of I^* , the generality of the model, and the statistical hardness of the task as well. From an experimental standpoint, tuning Q modifies the statistical complexity of the scene, and makes it possible to control the coarseness of the segmentation, with the possibility to build a hierarchy of coarse-to-fine (multiscale) segmentations of an image [3].

3 THEORETICAL ANALYSIS AND ALGORITHMS

For the sake of simplicity, we first state our theoretical results for a single color band (e.g., gray-level). On this basis, the extension of the results to more numerical channels, such as **RGB**, does not require an involved analysis: It is presented in Section 3.3. Recall that it is enough to give a merging predicate and an order to test region mergings, to completely define our segmentation algorithm.

3.1 Merging Predicate

Our first result is based on the following theorem.

Theorem 1 (The independent bounded difference inequality, [11]). *Let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be a family of n independent r.v. with X_k taking values in a set A_k for each k . Suppose that the real-valued function f defined on $\prod_k A_k$ satisfies $|f(\mathbf{x}) - f(\mathbf{x}')| \leq c_k$ whenever vectors \mathbf{x} and \mathbf{x}' differ only in the k th coordinate. Let μ be the expected value of the r.v. $f(\mathbf{X})$. Then, for any $\tau \geq 0$,*

$$\Pr(f(\mathbf{X}) - \mu \geq \tau) \leq \exp\left(-2\tau^2 / \sum_k (c_k)^2\right). \quad (1)$$

From this theorem, we obtain the following result on the deviation of observed differences between regions of I . Here, the notation $\mathbf{E}(R)$ for some arbitrary region R is the expectation over all corresponding statistical pixels of I^* of their sum of expectations of their Q r.v. for the single color band, and \bar{R} is the observed average of this color band.

Corollary 1. *Consider a fixed couple (R, R') of regions of I . $\forall 0 < \delta \leq 1$, the probability is no more than δ that*

$$|(\bar{R} - \bar{R}') - \mathbf{E}(\bar{R} - \bar{R}')| \geq g\sqrt{\frac{1}{2Q} \left(\frac{1}{|R|} + \frac{1}{|R'|} \right) \ln \frac{2}{\delta}}. \quad (2)$$

Proof. Suppose we shift the value of the outcome of one r.v. among the $Q(|R| + |R'|)$ possible for the couple (R, R') . $|\bar{R} - \bar{R}'|$ is subject to a variation of at most $c_R = g/(Q|R|)$ when this modification affects region R (among $Q|R|$ possible), and at most $c_{R'} = g/(Q|R'|)$ for a change inside R' (among $Q|R'|$ possible). We get $\sum_k (c_k)^2 = Q(|R|(c_R)^2 + |R'|(c_{R'})^2) = (g^2/Q)((1/|R|) + (1/|R'|))$. Using the fact that

the deviation with the absolute value is at most twice that without, and using Theorem 1 (solving for τ) brings our result. \square

Suppose we do N merging tests in I . Then, with probability $\geq 1 - (N\delta)$, all couples of regions (R, R') whose merging is tested shall satisfy $|(\bar{R} - \bar{R}') - \mathbf{E}(\bar{R} - \bar{R}')| \leq b(R, R')$, with $b(R, R')$ the right member of Corollary 1. Remark that N is small: for a single-pass algorithm, $N < |I|^2$. In our 4-connectivity setting (each pixel is connected to its north, south, east, and west neighbors when they exist), we even have $N < 2|I|$. What we really need to test the merging of two observed regions R and R' is a predicate accurate enough when the pixels of $R \cup R'$ come from the *same* statistical region of I^* . From this standpoint, using Corollary 1 to devise a merging predicate is straightforward: In this case, we have $\mathbf{E}(\bar{R} - \bar{R}') = 0$ and, thus, with high probability, the deviation $|\bar{R} - \bar{R}'|$ does not exceed $b(R, R')$. The merging predicate on two candidate regions R and R' could thus be “merge R and R' iff $|\bar{R} - \bar{R}'| \leq b(R, R')$,” with $b(R, R')$ a *merging threshold*. We shall see hereafter that such a predicate is *optimistic*: Under some assumption, it tends sometimes to favor overmerging (i.e., it does more merges than necessary to actually recover I^*), but this phenomenon formally remains quantitatively small. For both theoretical and practical considerations, we are going to replace this merging predicate by one slightly more optimistic, i.e., with a larger merging threshold. This one turns out to theoretically incur the same error (up to low order terms), and it gives very good visual results. Let \mathcal{R}_l be the set of regions with l pixels and $b(R) = g\sqrt{(1/(2Q|R|)) \ln(|\mathcal{R}_l|/\delta)}$. Remark that provided regions R and R' are not empty,

$$b(R, R') \leq \sqrt{b^2(R) + b^2(R')} < b(R) + b(R'). \quad (3)$$

Hereafter, we prove a quantitative bound on the error obtained with the largest quantity (the right one) used as merging threshold: it holds for both others as well. The center quantity is the merging threshold we use. An upperbound on $|\mathcal{R}_l|$ makes it quite reasonable with regard to $b(R, R')$. Considering that a region is an unordered bag of pixels (each color channel is given 0, 1, ..., l pixels), we may fix $|\mathcal{R}_l| \leq (l+1)^{\min\{l, g\}}$ (we have $l+1$ choices for the number of pixels having each color channel, which makes $|\mathcal{R}_l| \leq (l+1)^g$, and then we reduce this large upperbound by counting the duplicates for $l < g$). To summarize, our merging predicate is:

$$\mathcal{P}(R, R') = \begin{cases} \text{true} & \text{if } |\bar{R}' - \bar{R}| \leq \sqrt{b^2(R) + b^2(R')} \\ \text{false} & \text{otherwise.} \end{cases} \quad (4)$$

3.2 Order in Merging

The order in which we test the merging of regions follows a simple invariant \mathcal{A} which we define as follows:

- $(\mathcal{A}) \stackrel{\text{def}}{=} \text{when any test between two (parts of) true regions occurs, that means that all tests inside each of the two true regions have previously occurred.}$

It is crucial to note that \mathcal{A} does not postulate the knowledge of the segmentation of I^* . To make it clear why we should strive to fulfill \mathcal{A} , let us first recall the three types of error a segmentation can suffer. First, undermerging represents the case where one or more regions obtained are strict subparts of true regions. Second, overmerging represents the case where some regions obtained strictly contain more than one true region. Third, there is the “hybrid” (and most probable) case where some regions obtained contain more than one strict subpart of true regions. We have already partially outlined this in the preceding section related to the merging predicate: together with \mathcal{P} (4), \mathcal{A} makes it possible to control the segmentation error from both the qualitative and quantitative standpoints. The next theorem states that only overmerging occurs with high probability. In this theorem, we define $s^*(I)$ as the set of regions of the ideal (optimal) segmentation of I (defined from I^* , see Fig. 3) and $s(I)$ the set of regions in our segmentation of I .

Theorem 2. *With probability $\geq 1 - \mathcal{O}(|I|\delta)$, the segmentation on I satisfying \mathcal{A} is an overmerging of I^* , that is: $\forall O \in s^*(I), \exists R \in s(I) : O \subseteq R$.*

Proof. From Corollary 1, with probability $> 1 - (N\delta) = 1 - \mathcal{O}(|I|\delta)$, any couple of regions (R, R') coming from the same statistical region of I^* , and whose merging is tested, satisfy $|\bar{R} - \bar{R}'| \leq b(R, R')$. Since $b(R, R') \leq \sqrt{b^2(\bar{R}) + b^2(\bar{R}'')}$, our merging predicate $\mathcal{P}(R, R')$ (4) would authorize the merging of R and R' . Using the fact that \mathcal{A} holds together with this property, we first rebuild all true regions of I^* , and then eventually make some more merges: The segmentation obtained is an overmerging of I^* with high probability, as claimed. \square

The next theorem shows a quantitative upperbound on the error incurred with respect to the optimal segmentation. We define this error as the weighted average of the (absolute) channel differences over all nonempty intersections of regions between $s^*(I)$ and $s(I)$:

$$Err(s(I)) = \mathbf{E}_{R \cap O, R \in s(I), O \in s^*(I)} |\mathbf{E}(O) - \mathbf{E}(R)|, \quad (5)$$

with \mathbf{E} (slanted) denoting the expectation with associated probability measure $\mu(R \cap O) = |R \cap O|/|I|$.

Theorem 3. $\forall 0 < \delta < 1$, with probability $\geq 1 - \mathcal{O}(|I|\delta)$:

$$Err(s(I)) \leq \mathcal{O}\left(g \sqrt{\frac{|s^*(I)| \ln |s^*(I)|}{|I|Q}} \left(\ln \frac{1}{\delta} + g \ln |I|\right)\right). \quad (6)$$

(Proof omitted.) This theorem is interesting for three (mostly) theoretical reasons. First, the constant hidden in the big-Oh notation is small ($< \sqrt{6}$); second, it is proven for the largest merging threshold in (3). Last, if we ignore log-terms, the error incurred by our segmentation is driven by $g \sqrt{|s^*(I)|}/(|I|Q)$, a close order approximation to the optimum [12].

3.3 Color Images

The merging predicate for the RGB setting is:

$$\mathcal{P}(R, R') = \begin{cases} \text{true} & \text{if } \forall a \in \{\mathbf{R}, \mathbf{G}, \mathbf{B}\}, \\ & |\bar{R}'_a - \bar{R}_a| \leq \sqrt{b^2(\bar{R}) + b^2(\bar{R}')} \\ \text{false} & \text{otherwise.} \end{cases} \quad (7)$$

Here, \bar{R}_a denotes the observed average for color channel a in region R . Provided invariant \mathcal{A} holds as in Section 3.2, our predicate preserves overmerging, and the same bound as that of Theorem 3 holds on the error if we measure it as the sum of errors over the three color channels.

3.4 Our Algorithm: SRM

In 4-connectivity, there are $N < 2|I|$ couples of adjacent pixels. Let S_I be the set of these couples. Let $f(p, p')$ be a real-valued function, with p and p' pixels of I . Our segmentation algorithm, SRM (for Statistical Region Merging) is simple. We first sort the couples of S_I in increasing order of $f(\cdot, \cdot)$, and then traverse this order *only once*. We make for any current couple of pixels $(p, p') \in S_I$ for which $R(p) \neq R(p')$ ($R(p)$ stands for the current region to which p belongs) the test $\mathcal{P}(R(p), R(p'))$, and merge $R(p)$ and $R(p')$ iff it returns true. The objective is obviously to choose $f(\cdot, \cdot)$ so as to approximate \mathcal{A} as best as possible.

The next section reviews some choices we have made for $f(\cdot, \cdot)$, each of constant time computation. Because we do not update the list of merging tests after merging two regions, a simple ordering based on radix sorting with color differences as the keys yields a preordering time complexity $\mathcal{O}(|I| \log(g))$ —linear in $|I|$ —for our basic implementations of SRM. The merging steps afterward are space/time computational optimal [13], which makes SRM also optimal from both standpoints. The execution time of our basic implementation of SRM, which is not optimized, segments our largest images (512×512) in about one second on an Intel Pentium® IV 2.40 GHz processor.

4 EXPERIMENTAL RESULTS

Our model of image generation makes implicitly the assumption that observed color variations inside true regions should reasonably be smaller than between true regions. Such an assumption is made explicit in [8]. Thus, a good way to approximate \mathcal{A} is to capture the between-pixel local gradients, and then compute their maximal per-channel variation in $f(\cdot, \cdot)$: $f(p, p') = \max_{a \in \{\mathbf{R}, \mathbf{G}, \mathbf{B}\}} f_a(p, p')$. Below, we review some experiments using SRM. The reader may keep in mind that, unless otherwise stated, the values of the parameters of SRM are the same for all images: $\delta = 1/(6|I|^2)$ (Corollary 1) and $Q = 32$. Furthermore, the images are used as they are, i.e., without any preprocessing. Therefore, there is no extensive domain nor image-dependent tuning of the parameters.

4.1 Basic Choices for $f(\cdot, \cdot)$

We have tested two basic choices to compute $f_a(p, p')$. The simplest choice is to pick directly the pixel channel values (p_a and p'_a):

$$f_a(p, p') = |p_a - p'_a|. \quad (8)$$

Our second choice for $f_a(\cdot, \cdot)$ consists of extending convolution kernels classically used in edge detection for pixel-wise gradient estimation. In 4-connectivity, neighbor pixels are either horizontal or vertical. Thus, we only need $\hat{\partial}_x$ or $\hat{\partial}_y$ between neighbor pixels p and p' , for each color channel. We have chosen Sobel convolution filters, where smoothing is performed by the convolution mask $\begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$ and the derivative filter is performed by the convolution mask $\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$.




Fig. 4. An experimental display of the importance of sorting. Regions obtained in the segmentations are gray-level averaged with white borders.

Regardless of the choice of $f_a(.,.)$, the fact that we do not reorder the merging list during the merging steps might appear to be quite a strong constraint for efficient approximations to \mathcal{A} . The following simple experiment is an advocacy that it is not the case, and it uses our simplest implementation of $f_a(.,.)$ (8). Fig. 4 displays the result of SRM on gray-level images (with $\delta = 1/|I|^2$), and the result of the same algorithm in which the order is replaced by a conventional scanning of the image (from left to right and top to bottom) [13]. The preordering clearly manages dramatic improvements over conventional scanning.

Fig. 5 presents experiments obtained with our two methods for computing $f_a(.,.)$ on images for which results are visually different. On images with significant color




Fig. 5. Comparison of SRM without specific gradient estimation ((8), center images) and with convolution kernels (right images). Regions are color averaged with white borders.




Fig. 6. Sample results comparing both versions of SRM and [8]. Segmentation conventions are [8]s: region colors are chosen at random.

gradients, there is a visual advantage to SRM(w) (e.g., cup). Notice also the segmentation of SRM(w) on bldg, a picture exhibiting a large amount of motion blur. Remark from squirrel that SRM is able to isolate regions with high variability (e.g., the grass), and obtains results even better than [9] on the squirrel image: Their segmentation, although tailor-made for textured images, obtains a segmentation of the grass with many holes, a common drawback of region-merging techniques [9].

4.2 Random Noise Corruption

We have chosen to study the way SRM handles noise with our two choices for $f_a(.,.)$, against two hard noise types. Each color channel $\in \{\mathbf{R}, \mathbf{G}, \mathbf{B}\}$ of each pixel $\in I$ is transformed with probability $q \in [0, 1]$ into a new value:

- chosen uniformly in $\{1, 2, \dots, g\}$ for transmission noise ($t(q)$), or
- chosen uniformly in $\{1, g\}$ (the extremes) for salt and pepper noise ($s(q)$).

Fig. 6 shows different images corrupted with increasing amounts of noise, and the results of [8] and SRM. From 45 percent noise, the results of [8] appear to be random, whereas SRM still manages to find most interesting regions of the images. However, on some images, SRM obtained a brutal degradation of its performances for significant noise levels and for both ways of computing $f_a(.,.)$, indicating that the algorithm seems to reach its limits. To handle larger noise




Fig. 7. Results and comparisons with related approaches of SRM with our noise extensions for (8) (w/o) and Sobel-type filters (w), with $\Delta = 10$. (See text for the segmentation conventions.)

corruption, we have extended both solutions for $f_a(\cdot, \cdot)$ to more robust estimations, paying no significant additional computational cost. First, we replace p_a (8), by a moving average over a region defining a neighborhood around the pixel: $f_a(p, p') = |\overline{N_p(p')}_a - N_{p'}(p)_a|$. Here, $N_p(p')$ is the region defined by the set of points of I that are within Manhattan distance $\leq 2\Delta$ to p' (for some integer $\Delta \geq 0$), and that are closer to p' than they are to p . Whenever $\Delta = 0$, this expression matches that of (8). Second, we replace our convolution filter by larger ones, where smoothing is performed by the convolution mask $[1 \ 2 \ \dots \ \Delta + 1 \ \Delta \ \dots \ 1]$, and the derivative filter on each color channel is replaced by a local least-square linear regression on points whose abscissae are those of the convolution mask $[-\Delta \ -\Delta + 1 \ \dots \ \Delta]$, and ordinates defined by color channels of the corresponding pixels. Whenever $\Delta = 1$, this matches our Sobel filter's extension in Section 4.1.

Using radix sorting again with $f(\cdot, \cdot)$ values as the keys, the whole time complexity of our modifications of SRM becomes $\mathcal{O}(|I|(\log g + \Delta))$, which is still linear in $|I|$ if Δ is constant.

Fig. 7 reports results on the `castle` of Fig. 1. Conventions for the segmentations results are as follows: [10]'s regions are averaged with the original colors, [8]'s are averaged with random colors, and SRMs follow [10]'s (with white bordered regions). Notice that the number of regions found by [10], [8] explode with corruption, a phenomenon which does not appear for SRM modified. The segmentation time for the three algorithms gives a clear advantage to [8] and SRM modified. This image gives a slight advantage to SRM(w/o) over SRM(w) for noise handling, but we have remarked that both versions perform quite similarly on average. The best value of Δ , which controls the local number of pixels on which each gradient approximation is




Fig. 8. Results on images with occlusions. The largest regions found by SRM are displayed for each image.

computed, seems to vary between images, but it always has to remain small so as not to obtain "boxy-shaped" regions. In this respect, $\Delta = 10$ is not far from the maximal value.

4.3 Handling Occlusions

In our model, occlusions make it necessary to relax the 4-connexity constraint on the statistical regions of I^* . Handling them from an experimental point of view is simple: We first run SRM as already presented. Then, in a second stage, we run it again with two major modifications on the preordering step: We replace the pixels of I by the regions found after the first step, and replace the 4-adjacency graph between pixels by the clique graph between regions. We also replace $f_a(\cdot, \cdot)$ in (8) by: $f_a(R, R') = |\overline{R'}_a - \overline{R}_a|$. Radix sorting with the $f(\cdot, \cdot)$ values as the keys brings an overall time complexity $\mathcal{O}((|I| + k^2) \log g)$, where k is the number of regions found after the first step. The fact that our approach relies on slight overmerging tends to narrow the influence of k in the complexity. Fig. 8 shows some results obtained, on which SRM has managed to rebuild accurately the principal occluded regions (such as the road on the road image, despite the relative noise of this video-extracted picture).

4.4 Controlling the Scale of the Segmentation

Some authors have emphasized the need to control the coarseness of a segmentation [7], [3], [4]. The objective of multiscale segmentation is to get a hierarchy of segmentations at different scales, and get at each scale a level of details compatible with the perceptual organization of the image at this scale. In our case, controlling the scale is made easy with the tuning of parameter Q : The smaller it is, the harder is the statistical estimation task, and the less




Fig. 9. Segmentations of SRM on image `house`, for different values of Q . Regions found are white-bordered (see text for details).

numerous are the regions in the final segmentation. To visualize this, we run SRM with $\Delta = 2$ for our extension of Sobel convolution filters (see Section 4.2), and making Q range through the values 1, 2, 4, \dots , 256. Fig. 9 presents the results obtained on `house` image. It is interesting to note that as Q increases, the regions found are getting smaller, but they often correspond to smaller perceptual regions of the image at different scales (e.g., the house gets segmented gradually, from itself as a whole until all its details are gradually extracted: facades, windows, rooftops, etc.).

5 CONCLUSION

In this paper, we propose a segmentation algorithm based on the idea that perceptual grouping with region merging has to catch the big picture of a scene by only having primary local glimpses on it. Our algorithm is based on a model of image generation which captures the idea that grouping is an inference problem. This provides us with a simple merging predicate, and a simple ordering in merges which, with high probability, both suffers only one source of error (overmerging), and achieves with high probability a low error in segmentation. It can be reliably approximated by very fast segmentation algorithm, SRM, which from our experiments tends indeed to satisfy our goal of image segmentation. Experiments display the ability of the approach to cope with significant noise corruption, handle occlusions, and perform scale-sensitive segmentations.

ACKNOWLEDGMENTS

The authors would like to thank the reviewers for their insightful comments on this paper. This work was done while Richard Nock was visiting Sony CSL Tokyo. Resources related to SRM may be obtained from the authors Web pages.

REFERENCES

- [1] D. Forsyth and J. Ponce, *Computer Vision—A Modern Approach*. Prentice Hall, 2003.
- [2] Z. Wu and R. Leahy, "An Optimal Graph Theoretic Approach to Data Clustering," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, pp. 1101-1113, 1993.
- [3] E. Sharon, A. Brandt, and R. Basri, "Fast Multiscale Image Segmentation," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 70-77, 2000.
- [4] E. Sharon, A. Brandt, and R. Basri, "Segmentation and Boundary Detection Using Multiscale Intensity Measurement," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 469-476, 2001.
- [5] J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, pp. 888-905, 2000.
- [6] A. Barbu and S.-C. Zhu, "Graph Partition by Swendsen-Wang Cuts," *Proc. Ninth IEEE Int'l Conf. Computer Vision*, pp. 320-327, 2003.
- [7] M. Galun, E. Sharon, R. Basri, and A. Brandt, "Texture Segmentation by Multiscale Aggregation of Filter Responses and Shape Elements," *Proc. Ninth IEEE Int'l Conf. Computer Vision*, pp. 716-725, 2003.
- [8] P.F. Felzenszwalb and D.P. Huttenlocher, "Image Segmentation Using Local Variations," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 98-104, 1998.
- [9] S.-C. Zhu and A. Yuille, "Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, pp. 884-900, 1996.
- [10] D. Comaniciu and P. Meer, "Robust Analysis of Feature Spaces: Color Image Segmentation," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 750-755, 1997.
- [11] C. McDiarmid, "Concentration," *Probabilistic Methods for Algorithmic Discrete Math.*, M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed, eds., pp. 1-54, Springer Verlag, 1998.
- [12] M.J. Kearns and Y. Mansour, "A Fast, Bottom-Up Decision Tree Pruning Algorithm with Near-Optimal Generalization," *Proc. 15th Int'l Conf. Machine Learning*, pp. 269-277, 1998.
- [13] C. Fiorio and J. Gustedt, "Two Linear Time Union-Find Strategies for Image Processing," *Theoretical Computer Science*, vol. 154, pp. 165-181, 1996.



learning, data mining, computational complexity, and image processing.



Frank Nielsen received the BSc and MSc degrees from Ecole Normale Supérieure (ENS) of Lyon (France) in 1992 and 1994, respectively. He defended his PhD thesis on adaptive computational geometry prepared at INRIA Sophia-Antipolis (France) under the supervision of Professor J.-D. Boissonnat in 1996. As a civil servant of the University of Nice (France), he gave lectures at the engineering schools ESSI and ISIA (Ecole des Mines). In 1997, he served in the army as a scientific member in the computer science laboratory of Ecole Polytechnique. In 1998, he joined Sony Computer Science Laboratories Inc., Tokyo (Japan), as a researcher. His current research interests include geometry, vision, graphics, learning, and optimization.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.