Multi-labelled Image Segmentation in Irregular, Weighted Networks: A Spatial Autocorrelation Approach

Raphaël Ceré and François Bavaud
Department of Geography and Sustainability, University of Lausanne, Switzerland

Keywords: Free Energy, Image Segmentation, Iterative Clustering, K-means, Laplacian, Modularity, Multivariate Features, Ncut, Soft Membership, Spatial Autocorrelation, Spatial Clustering.

Abstract: Image segmentation and spatial clustering both face the same primary problem, namely to gather together spatial entities which are both spatially close and similar regarding their features. The parallelism is particularly obvious in the case of irregular, weighted networks, where methods borrowed from spatial analysis and general data analysis (soft K-means) may serve at segmenting images, as illustrated on four examples. Our semi-supervised approach considers soft memberships (fuzzy clustering) and attempts to minimize a free energy functional made of three ingredients : a within-cluster features dispersion (hard K-means), a network partitioning objective (such as the Ncut or the modularity) and a regularizing entropic term, enabling an iterative computation of the locally optimal soft clusters. In particular, the second functional enjoys many possible formulations, arguably helpful in unifying various conceptualizations of space through the probabilistic selection of pairs of neighbours, as well as their relation to spatial autocorrelation (Moran's I).

1 INTRODUCTION

Regional data analysis, as performed on geographic information systems, deals with a notion of "*where*" (the spatial disposition of regions), a notion of "*what*" (the regional features), and a notion of "*how much*" (the relative importance of regions, as given by their surface or the population size). The data define a *marked, weighted network*, generally *irregular* (think e.g. of administrative units): weighted vertices represent the regions, weighted edges measure the proximity between regions, on which uni- or multivariate features (the marks) are defined.

Much the same can be said of an image made of pixels, that is a collection of elements embedded in a bidimensional layout. The regularity of the setup (regular grid, uniform weights, binary regular adjacencies) is exploited in most segmentation algorithms, but the latter may become inadapted, precisely, under irregular situations, such as pixels of various sizes or importance, aggregated pixels, irregular boundaries or connectivites, multi-layered or partially missing data.

This paper proposes an iterative algorithm for semi-supervised image segmentation, directly adapted from regional clustering procedures in spatial analysis, themselves originating in (non-marked) network clustering. In a nutshell, the purely spatial standard procedures aimed at *Ncut minimisation* (Shi and Malik, 2000; Grady and Schwartz, 2006) or *modularity maximization* (Newman, 2006) are enriched with a features dissimilarity term, central in the *K-means approach*, and further regularized by an *entropy term*, favoring the emergence of soft clusters, and allowing a iterative computation of locally optimal solutions.

The fields of spatial analysis, in particular spatial clustering, on one hand, and image segmentation on the other hand, seem currently to be investigated by distinct, non-overlapping communities. Yet, both communities arguably share the same *what-versus-where-trading* challenge, aimed at obtaining clusters both homogeneous and connected.

In this work in progress, we first define the main quantities of interest and the iterative algorithm itself (section 2). The definition of the canonical measure of *image homogeneity*, namely *Moran's I* in the present spatial context, is recalled in section 2.3. Section (3) illustrates the method on four case sudies. A conclusion (section 4) lists some further research lines, to be investigated in a close future. The appendix details the construction of the so-called exchange matrix, as well as the test of spatial autocorrelation in a weighted setting.

Ceré, R. and Bavaud, F.
Multi-labelled Image Segmentation in Irregular, Weighted Networks: A Spatial Autocorrelation Approach.
DOI: 10.5220/0006322800620069
In *Proceedings of the 3rd International Conference on Geographical Information Systems Theory, Applications and Management (GISTAM 2017)*, pages 62-69
ISBN: 978-989-758-252-3

2 DEFINITIONS AND FORMALISM

The formalism we consider extends the *spatial autocorrelation formalism* used in Quantitative Geography and Spatial Econometrics to the case of weighted, irregular regions, as well as to multivariate features. It turns out to be broad enough to provide a flexible framework for unsupervised or semi-supervised *generalized image* segmentation, where the "generalized images" under consideration can be made of irregular pixels, irregularly inter-connected, and endowed with multivariate numerical features.

Space as a weighted network: the exchange matrix E

Specifically, consider n regions (generalized pixels) with relative weights $f_i > 0$, summing to $f_\bullet = \sum_{i=1}^n f_i = 1$, together with an $n \times n$ symmetric non-negative *exchange matrix* $E = (e_{ij})$, and *weight-compatible* in the sense $e_{i\bullet} = \sum_{j=1}^n e_{ij} = f_i$. The exchange matrix E interprets as a joint probability $p(i,j) = e_{ij}$ to select the pair of regions i and j (edges), and defines a weighted unoriented network. Its margins interpret as the probability $p(i) = f_i$ to select region i (vertices).

Weight-compatible exchange matrices E define a continuous neighborhood relation between regions. They can be constructed from f and the adjacency matrix A, or from another spatial proximity of distance matrix (see the appendix). The row-standardized matrix of spatial weights $W = (w_{ij})$ of spatial autoregressive models obtains as $w_{ij} = e_{ij}/f_i$, and constitutes the transition matrix of a Markov chain with stationary distribution f.

Multivariate features: the dissimilarity matrix D

Regional features can consist of univariate grey levels, multivariate color or spectral intensities, or (in a geographical context) any regional variable such as the proportions of specific land uses, population density, proportion of party B voters, etc. Multivariate characteristics x_i are suitably combined into $n \times n$ squared Euclidean dissimilarities $D_{ij} = \|x_i - x_j\|^2$.

2.1 Image Segmentation (Regional Clustering)

A soft *regional clustering* or *image segmentation* into m groups is described by non-negative $n \times m$ *membership matrix* $Z = (z_{ig})$ with $z_{ig} = p(g|i)$ denotes the probability that region (pixel) i belongs to group g, and obeys $z_{i\bullet} = \sum_{g=1}^m z_{ig} = 1$. The weights ρ of the corresponding groups obtain as $\rho_g = \sum_i f_i z_{ig}$, and the regional distribution of groups as $f_i^g = f_i z_{ig}/\rho_g = p(g|i)$, obeying $f_\bullet^g = 1$. The region-group dependency can be measured by the *mutual information*

$$\mathcal{K}[Z] = \sum_{ig} p(i,g) \ln \frac{p(i,g)}{p(i)p(g)} = \sum_{ig} f_i z_{ig} \ln \frac{z_{ig}}{\rho_g} \quad (1)$$

A good clustering should consist of homogeneous groups made of regions not too dissimilar regarding their features, that is insuring a low value of the *within-group inertia* (Bavaud, 2009)

$$\Delta_W[Z] = \sum_{g=1}^m \rho_g \Delta_g \quad \text{where} \quad \Delta_g = \frac{1}{2} \sum_{ij} f_i^g f_j^g D_{ij} \quad (2)$$

A good clustering should also avoid to separate a pair of spatially strongly connected regions, that is to insure a low value of the *generalized cut*

$$\mathcal{G}[Z] = \frac{1}{2} \sum_g G(\rho_g) \sum_{ij} e_{ij} (z_{ig} - z_{jg})^2 \quad (3)$$

where $G(\rho) \geq 0$ is non-increasing in ρ. The choice $G(\rho) = 1/\rho$ amounts to the N-cut (Shi and Malik, 2000), while the choice $G(\rho) = 1$, *we shall adopt here*, is equivalent to the modularity criterium (Newman, 2006).

We consider the *regularized clustering problem*, consisting in finding a clustering Z minimizing the *free energy* functional

$$\mathcal{F}[Z] = \beta \Delta_W[Z] + \frac{\alpha}{2} \mathcal{G}[Z] + \mathcal{K}[Z] \quad (4)$$

The terms Δ_W, respectively $\mathcal{G}$, behaves as a features dissimilarity energy, respectively a spatial energy, favoring *hard* partitions obeying $z_{ig} = 0$ or $z_{ig} = 1$. By contrast, the regularizing entropy term $\mathcal{K}$ favors the emergence of soft clusterings. Setting $\alpha = 0$ yields the soft K-means algorithm (Gaussian mixtures), where β is an *inverse (dissimilarity) temperature*. Canceling the first-order derivative of the free energy with respect to z_{ig} under the constraints $z_{i\bullet} = 1$ yields the minimization condition

$$z_{ig} = \frac{\rho_g \exp(-\beta D_i^g - \alpha(\mathcal{L}z^g)_i)}{\sum_h \rho_h \exp(-\beta D_i^h - \alpha(\mathcal{L}z^h)_i)} \quad (5)$$

where $D_i^g = \sum_j f_i^g D_{ij} - \Delta_g$ is the squared Euclidean dissimilarity from i to the centroid of group g, and $(\mathcal{L}z^g)_i = z_{ig} - \sum_k w_{ik} z_{kg}$ is the *Laplacian* of z^g at i, comparing its value to the average value of its neighbourgs.

2.2 Iterative Procedure

Equation (5) can be solved iteratively, updating ρ_g, D_{ig} and $\mathcal{L}z^g$ at each step, and converges towards a *local minimum* Z^∞ of $\mathcal{F}[Z]$, which constitutes the searched for soft spatial partition or image segmentation. Z^∞ can be further hardened by assigning i to group $g = \arg\max_h z^\infty_{ih}$.

A semi-supervised implementation of the procedure, imposing the membership of a few pixels (and possibly breaking down the monotonic decrease of $\mathcal{F}[Z]$: see figure 5) goes as follow: first, the set Ω of the n regions is partitioned into two disjoint, non-empty sets, namely the user-defined *tagged regions* T, and the *free regions* F, with $\Omega = T \cup F$ and $T \cap F = \emptyset$. The tagged set T itself consists of m non-empty disjoint subregions $T = \cup_{\tau=1}^m T_\tau$ initially tagged with m distinct *strokes* applied on a small number of pixels: they form the *seeds* of the $g = 1, \ldots, m$ figures to be extracted, while the remaining regions will be assigned to the background numbered $g = 0$.

Memberships $Z = (z_{ig})$ consist of $n \times (m+1)$ non-negative matrices obeying $\sum_{g=0}^m z_{ig} = 1$. Their initial value Z^0 is set as

$$z^0_{ig} \quad \begin{cases} = 1 & \text{if } i \in F \text{ and } g = 0 \\ = 1 & \text{if } i \in T_\tau \text{ and } g = \tau \\ = 0 & \text{otherwise .} \end{cases} \qquad (6)$$

Iteration (5) is then performed. At the end of each loop, the tagged regions are reset to their initial values $z^0_{i\tau} = 1$ for all $i \in T_\tau$. After convergence, one expects the hardened clusters obtained by assigning i to group $g = \arg\max_{h=0}^m z^\infty_{ih}$ to consist of m connected figures $g = 1, \ldots m$ each containing the tagged set T_g, as well as a remaining background supported on F.

The iterative image segmentation algorithm summarized below (table 1) requires

1) a vector of n weights $f_i > 0$ associated to each pixel or region
2) a vector of n grey levels or multivariate characteristics x_i
3) a $n \times n$ binary, symmetric, off-diagonal adjacency matrix A .

Retaining the soft memberships $Z^{(\infty)}$ after convergence could possibly provide a novel *edge detection mechanism* (figure 5(b)), to be investigated in a subsequent work. Namely, one expects $z^{(\infty)}_{ig} = 0$ or $z^{(\infty)}_{ig} = 1$ for pixels i located within the interior of homogeneous groups, while $0 < z^{(\infty)}_{ig} < 1$ for pixels i located at the boundary of two of more groups. As a consequence, the value of the *entropy* $H_i = -\sum_g z^{(\infty)}_{ig} \ln z^{(\infty)}_{ig}$ is presumably large for i located at the group frontiers.

Table 1: Summary of the iterative algorithm.

Begin
Compute the weight vector f ($f_i = 1/n$ for regular grids)
Compute the binary adjacency matrix A
For given $t > 0$, compute the weight-compatible exchange matrix $E(f,A,t)$ by (9), and the matrix of spatial weights as $w_{ij} = e_{ij}/f_i$
Compute the features dissimilarity matrix $D_{ij} = \|x_i - x_j\|^2$
Initialize the $n \times (m+1)$ membership matrix Z^0 as:
$z^0_{ig} = 1$ if $i \in F$ and $g = 0$
$z^0_{ig)} = 1$ if $i \in T_\tau$ and $g = \tau$
$z^0_{ig} = 0$ otherwise.
Loop : $Z^{(r+1)}$ for the r-th iteration, stop after convergence
Group weight : $\rho_g := \sum_i f_i z_{ig}$
Emission probabilities $f^g_i := \frac{f_i z_{ig}}{\rho_g}$
Dissimilarity to the centroid : $D^g_i := \sum_j f^g_j D_{ij} - \Delta_g$
Laplacian: $(\mathcal{L}z^g)_i = z_{ig} - \sum_j w_{ij} z_{jg}$
Compute $z^{(r+1)}_{ig}$ by (5)
Re-initialize $z^{(r+1)}_{ig} = \delta_{g\tau}$ for $i \in T_\tau$
Attribute $i \in F$ to the group $g = \arg\max_{h=0}^m z^{(\infty)}_{ih}$
End

Deterministic profiling

The above iterative algorithm has been developed with *Python* 2.7.12 (van Rossum, 1995) and performed on a *CPU Intel Core i7* two Core with a frequency 3.1GHz (Mac OS X 10.10.5). Figure 1 summarizes the "plain" computational performances (i.e. without further optimizing such as Parallel computing or Analysis of complexity) for the main ingredients at stake, as a function of the number of pixels n in a regular setting.

2.3 Spatial Autocorrelation: Moran's *I*

Average multivariate dissimilarities between regions are expressed by *inertias*, generalizing the univariate *variances*. The inertia between randomly selected regions, and the *local inertia* between neighbours, are respectively defined as

$$\Delta := \frac{1}{2} \sum_{i,j=1}^n f_i f_j D_{ij} \qquad \Delta_{\text{loc}} := \frac{1}{2} \sum_{i,j=1}^n e_{ij} D_{ij} \qquad (7)$$

Comparing the global versus local inertias provides a multivariate generalization of *Moran's I*, namely,

$$I := \frac{\Delta - \Delta_{\text{loc}}}{\Delta} \qquad (8)$$

whose values range in $[-1, 1]$. A large positive I is expected for an image made of large patchs characterized with constant features, or at least varying smoothly on average (spatial continuity = positive autocorrelation). A large negative I characterizes an image whose pixel features are contrasted, opposite to

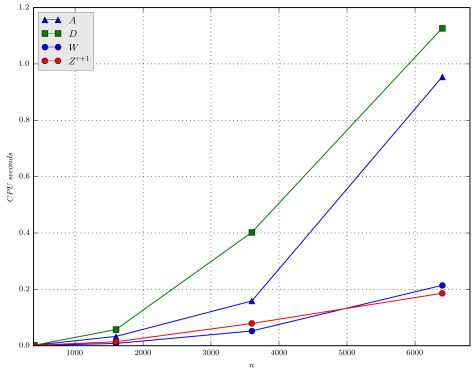

(a)

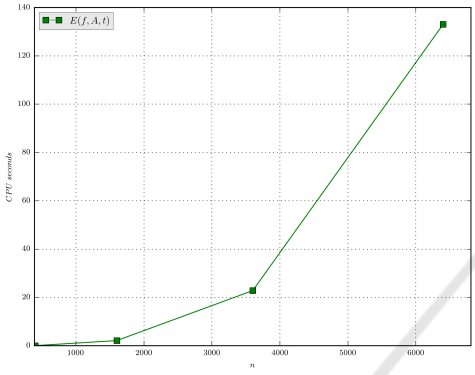

(b)

Figure 1: CPU times for a range of $n \in [400, 10000]$ pixels. (a) Computing of the binary "rook" adjacency matrix A (with the *Scikit-learn* package (Pedregosa et al., 2011)), the squared Euclidean profiles dissimilarity D, the Markov transition matrix W (for E given), as well as one single iteration $Z^{(r)} \to Z^{(r+1)}$ of the membership matrix. (b) The computation of the exchange matrix $E(f,A,t)$ by (9) involves a time-consuming complete eigen-decomposition.

their neighbours - such as a chess board with "rook" adjacency. See the appendix for testing the statistical significance of I.

3 ILLUSTRATIONS

We consider four small datasets:

- *Swiss votes*, an irregular, bivariate vector image defined by two votations results at the canton level (figure 2 and 3)
- *the Portrait*, a regular, trivariate (levels of red green blue) image (figure 4)
- *the Geometer*, a regular, univariate (grey levels) image (figure 5)
- *Thun*, a geographical univariate image depicting the building density at the hectometre scale (census blocks) in the region of Thun, Switzerland (figure 6).

We first study the segmentation procedure, before turing to spatial autocorrelation.

Figure 2 depicts the proportion of "yes" x_i, respectively y_i, for each canton i ($n = 26$) for two emblematic swiss popular initiatives, namely (a) the *Initiative against mass immigration* which was accepted February the 9th, 2014, on the tie (50.3% of the citizen, and the majority of cantons), respectively (b) the *Initiative for minimum wages* which was widely refused May the 25, 2014 (23.7% of the citizens, and the majority of cantons). Both results have been aggregated into the dissimilarity $D_{ij} = (x_i - x_j)^2 + (y_i - y_j)^2$, further rescaled in the range $[0,1]$. Non-uniform canton weights are defined relatively to the population 2015 density as $f_i = \text{POP}_i/\text{POP}_{total}$, and the "queen" adjacency scheme has been adopted.

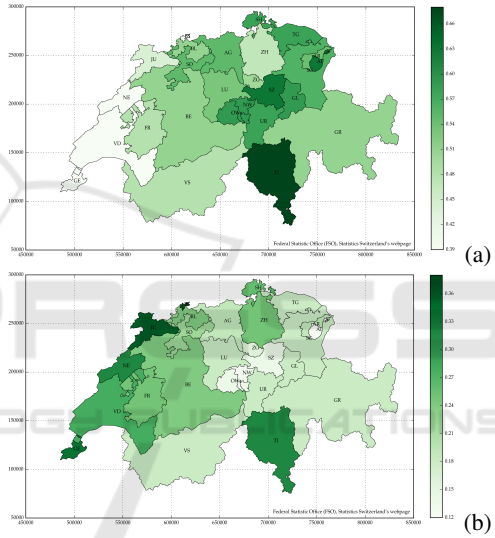

(a)

(b)

Figure 2: *Swiss votes:* percentage of "yes" for (a) the *Initiative against mass immigration* and for (b) the *Initiative for minimum wages*, at the canton level. ZH: Zürich, BE: Bern, LU: Luzern, UR: Uri, SZ: Schwyz, OW: Obwalden, NW: Nidwalden, GL: Glarus, ZG: Zug, FR: Fribourg/Freibourg, SO: Solothurn, BS: Basel-Stadt, BL: Basel-Landschaft, SH: Schaffhausen, AR: Appenzell-Aus., AI: Appenzell-Inn., SG: St. Gallen, GR: Grischun, AG: Aargau, TG: Thurgau, TI: Ticino, VD: Vaud, VS: Valais/Wallis, NE: Neuchâtel, GE: Genève, JU: Jura.

Figure 4 depicts a regular multivariate image (RGB) of size 50×54 with uniform weight vector $f_i = 1/n$ for $n = 2700$. Also, the binary adjacency matrix has been built under the "queen" scheme (8 neighbourghs), and D_{ij} is the sum of the squared differences between color intensities.

Segmentation evaluation

The *silhouette coefficient* $s \in [-1,1]$ is a measure of the tightness and separation associated to a dissimilarity-based clustering, and defines as $s = (b-a)/\max(a,b) \in [-1,1]$, where a is the mean intra-

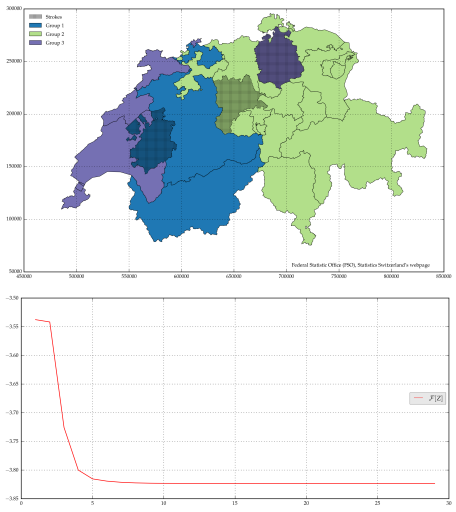

Figure 3: Swiss votes, continued: (a) hard assignment i to group g obtained from the intial strokes FR $\in$ Group 1, LU $\in$ Group 2 and ZH $\in$ Group 3 after 29 iterations, with $\beta = 1/0.007$ and $\alpha = 0.7143$. (b) depicts the decrease of the free energy functional $\mathcal{F}[Z]$ during the iteration.

cluster dissimilarity, and b the mean "nearest-cluster" dissimilarity (Rousseeuw, 1987).

Table 2 gives the silhouette coefficients s for three groups (including the background group) under four segmentation algorithms, possibly converted to hard clustering after convergence, namely: hard K-means with random initialization of centroids, EM Gaussian Mixture modeling with diagonal covariance matrix (Pedregosa et al., 2011), our own iterative clustering algorithm (table 1) with $\alpha = 0$ and β optimized so as to maximize s, and finally the iterative clustering algorithm (table 1) with the same value of β, but with $\alpha = 1$ in order to take into account the spatial contribution $\mathcal{L}z^g$. *Spatial autocorrelation*

Table 2: Segmentation evaluation : silhouette coefficients s under four datasets and four clustering algorithms (see text). The iterative algorithm with $\alpha = 0$ yields the optimal values $\beta_0 = 1/0.0052$ (Swiss votes), $\beta_0 = 1/0.0052$ (Portrait), $\beta_0 = 1/0.00342$ (Geometer) and $\beta_0 = 1/0.0052$ (Thun), and should be equivalent to the EM algorithm. The lower performance of the iterative algorithm (table 1) for the Swiss votes dataset is unexpected, and under current investigation.

Dataset	m+1	$s_{\text{K-means}}$	s_{EM}	$s_{\beta_0,\alpha=0}$	$s_{\beta_0,\alpha=1}$
Swiss votes	3	0.41	0.40	0.28	0.28
Portrait	3	0.57	0.57	0.56	0.56
Geometer	3	0.66	0.63	0.64	0.65
Thun	3	0.77	0.70	0.76	0.77

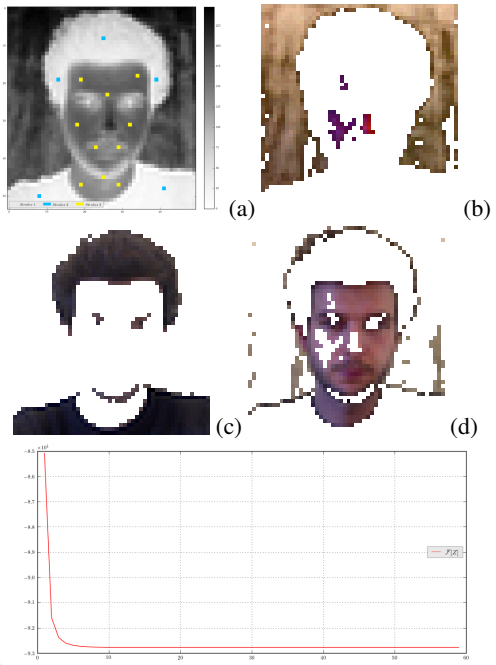

Figure 4: the Portrait: (a) initial configuration (intensity of red) and strokes (colored pixels); (b) hard assignment to background group $g = 0$; (c) hard assignment to $g = 1$, and (d) hard assignment to $g = 2$, obtained after 99 iterations, with $\beta = 1/0.001$ and $\alpha = 1$. Bottom: behavior of the free energy functional during the iteration.

Figure 7 presents the values of Moran's I (8), its expectation and variance under the null hypothesis of no autocorrelation, and its standardized z-value on which the normal test is based (see the appendix), for a diffusion time $t = 1$.

Figures 8, 9, 10 and 11 depict, for varying t, the values of Moran's I and its expectation $E_0(I)$ (top), its variance $\text{Var}_0(I)$ (middle) and its standardized z normal test value (bottom) for the four illustrations under consideration.

The freely adjustable parameter t controls the *neighbors range* of the diffusive Markov diffusion process $W(t)$, and the behaviour of the plotted quantities relatively to t should presumably reflect the size of patches consisting of similar pixels.

4 DISCUSSION

The present study constitutes a first attempt aimed to unify two little interacting domains – image segmentation and regional clustering – yet arguably very similar in their aims, and both relevant to spatial analysis. Spatial autocorrelation, in its general formulation (irregular, weighted networks) emerges as a common unifying paradigm, worth extending in the future

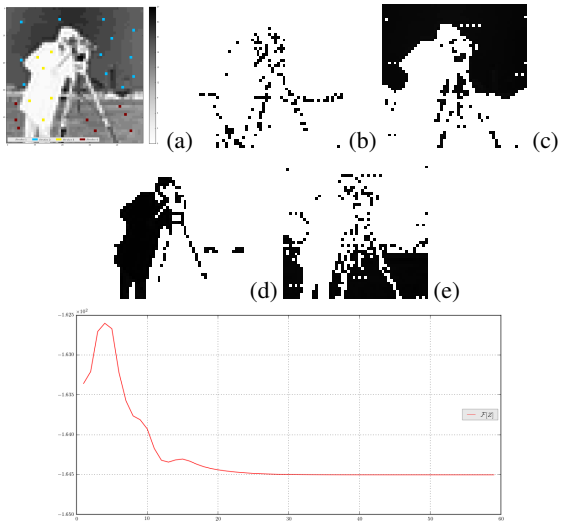

Figure 5: the Geometer: (a) raster univariate image (grey levels x_i range in the interval $[0,255]$) of size $n = 50 \times 50 = 2500$ with initial strokes; (b) background group $g = 0$; (c) $g = 1$; (d) $g = 2$, and (e) $g = 3$, obtained after 60 iterations, with parameters $\beta = 1/0.003$ and $\alpha = 0.017$. Bottom: behavior of the free energy functional during the iteration, initially *increasing*.

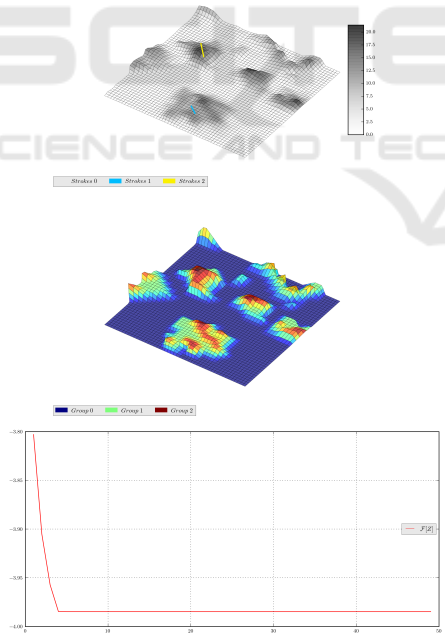

Figure 6: Thun example. Top: raster univariate image of size $n = 20 \times 24 = 480$, depicting the the population density per hectometer, rescaled in in the interval of $[0,27]$. Middle: groups $g = 0$ (blue, background), $g = 1$ (green) and $g = 2$ (purple), obtained after 50 iterations, with parameters $\beta = 1/0.0003$ and $\alpha = 0.17$. Bottom: behavior of the free energy.

I	$E_0(I)$	$\mathrm{Var}_0(I)$	z	Image(s)
0.65	0.33	0.0064	4.01	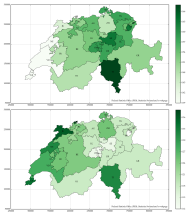
0.23	0.10	2.15×10^{-6}	88.92	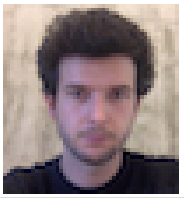
0.21	0.099	2.12×10^{-6}	74.41	
0.69	0.40	0.0001	24.49	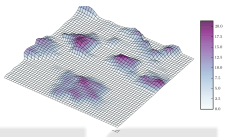

Figure 7: Moran's I and standardized test value z for the four illustrations.

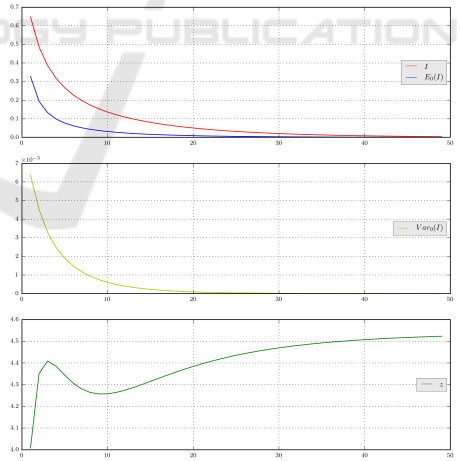

Figure 8: Swiss votes: Moran's I (top) and standardized z normal test value (bottom) for $t \in [1,50]$.

to local indicators (LISA) (Anselin, 1995). The proposed iterative approach can be considered as a generalization of the EM soft clustering for mixture models, containing an additional spatial term.

Alternative spatial or "Dirichlet-like" functionals $\mathcal{G}[Z]$, such as the N-cut specification mentioned in section (2), or "p-Laplacians" (Bougleux et al.,

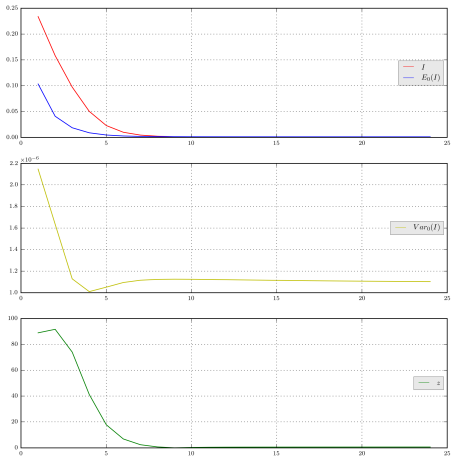

Figure 9: the Portait: Moran's I (top) and standardized z normal test value (bottom) for $t \in [1,25]$.

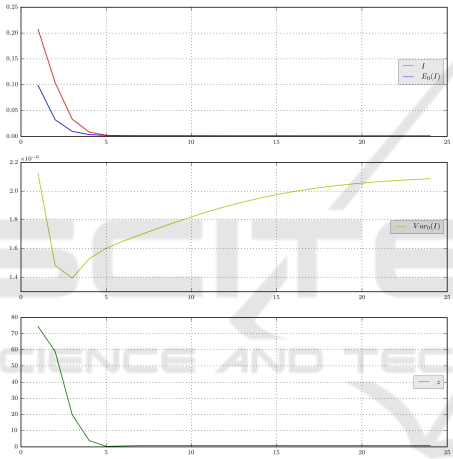

Figure 10: the Geometer: Moran's I (top) and standardized z normal test value (bottom) for $t \in [1,25]$.

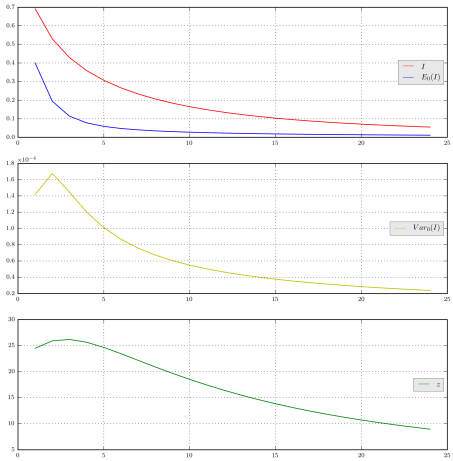

Figure 11: Thun: Moran's I (top) and standardized z normal test value (bottom) for $t \in [1,25]$.

2009) should be considered, as is the relation to the Mumford-Shah functional (Chambolle et al., 2012) (Bar et al., 2011). Note that $\mathcal{G}[Z]$ can be adjusted in two obvious, independent ways, namely by considering alternative weight-compatible specifications to $E(A,f)$, and by considering alternative forms $G(\rho)$ in (3).

Besides those energy-based approaches, the entropic part $\mathcal{K}[Z]$ can be introduced as a simple regularization term, or, as in the EM soft clustering, as originating from maximizing the log-likelihood of a statistical model (MAP approach), as in (Besag, 1986) or (Greig et al., 1989); see also (Couprie et al., 2011). Probabilistic formulations of the image segmentation problem should, as attested from spatial autocorrelation studies, benefit from results and experience in model-based clustering and latent modeling.

Alternatively, the virtue of random walk approaches, intimately connected to Markov chains and weighted networks, electric of not, has been attested in semi-supervised clustering in general, and image segmentation in particular (Grady, 2006). One key quantity is the probability that a random walk starting from a free pixel reaches first each of the m tagged regions – a problem well-known to be related to the Dirichlet differential problem with suitable boundary conditions, and to the computation of electric potentials (Doyle and Snell, 1984).

Yet, in the probabilistic formulations of the multi-source, multi-target random walks (Guex, 2016), the *capacities* and the edge *resistances* turn out to be *freely, separately adjustable* (Bavaud and Guex, 2012). In the image segmentation context, it is tempting to identify the capacity contribution as a spatial term enabling transitions between neighbors, and the resistance contribution as a barrier preventing transitions between too dissimilar pixels; to which extent this line of research could prove itself innovative and efficient should be addressed in a near future.

REFERENCES

Anselin, L. (1995). Local indicators of spatial association-lisa. *Geographical analysis*, 27(2):93–115.

Bar, L., Chan, T. F., Chung, G., Jung, M., Kiryati, N., Mohieddine, R., Sochen, N., and Vese, L. A. (2011). Mumford and shah model and its applications to image segmentation and image restoration. In *Handbook of mathematical methods in imaging*, pages 1095–1157. Springer.

Bavaud, F. (2009). Aggregation invariance in general clustering approaches. *Advances in Data Analysis and Classification*, 3(3):205–225.

Bavaud, F. (2014). Spatial weights: constructing weight-compatible exchange matrices from proximity matri-

ces. In *International Conference on Geographic Information Science*, pages 81–96. Springer.

Bavaud, F. and Guex, G. (2012). Interpolating between random walks and shortest paths: a path functional approach. In *International Conference on Social Informatics*, pages 68–81. Springer.

Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 259–302.

Bougleux, S., Elmoataz, A., and Melkemi, M. (2009). Local and nonlocal discrete regularization on weighted graphs for image and mesh processing. *International journal of computer vision*, 84(2):220–236.

Chambolle, A., Cremers, D., and Pock, T. (2012). A convex approach to minimal partitions. *SIAM Journal on Imaging Sciences*, 5(4):1113–1158.

Couprie, C., Grady, L., Najman, L., and Talbot, H. (2011). Power watershed: A unifying graph-based optimization framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(7):1384–1399.

Doyle, P. G. and Snell, J. L. (1984). *Random walks and electric networks*. Mathematical Association of America,.

Grady, L. (2006). Random walks for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 28(11):1768–1783.

Grady, L. and Schwartz, E. L. (2006). Isoperimetric graph partitioning for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 28(3):469–475.

Greig, D. M., Porteous, B. T., and Seheult, A. H. (1989). Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279.

Guex, G. (2016). Interpolating between random walks and optimal transportation routes: Flow with multiple sources and targets. *Physica A: Statistical Mechanics and its Applications*, 450:264–277.

Newman, M. E. J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65.

Schneider, M. H. and Zenios, S. A. (1990). A comparative study of algorithms for matrix balancing. *Operations research*, 38(3):439–455.

Shi, J. and Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905.

van Rossum, G. (1995). Python tutorial. *Technical Report CS-R9526, Centrum voor Wiskunde en Informatica (CWI), Amsterdam*.

APPENDIX

Computing $E(A,f)$

Constructing an exchange matrix E (section 2) both weight-compatible (that is obeying $E\mathbf{1} = f$, where the regional weights f are given) and reflecting the spatial structure contained in the binary adjacency matrix $A = (a_{ij})$ is not trivial, nor that difficult either.

A natural attempt consists in determining a vector c such that $e_{ij} = c_i\, a_{ij}\, c_j$. The problem can be essentially solved by Sinkhorn iterative fitting (Schneider and Zenios, 1990)

In this paper, we alternatively consider A as the infinitesimal generator of a continuous Markov chain at time $t > 0$, which yields the *diffusive specification* (Bavaud, 2014)

$$E \equiv E(A,f,t) = \Pi^{1/2} \exp(-t\Psi)\, \Pi^{1/2} \qquad (9)$$

where $\Pi = \mathrm{diag}(f)$, and

$$\Psi = \Pi^{-1/2} \frac{LA}{\mathrm{trace}(LA)} \Pi^{-1/2} \qquad (LA)_{ij} = \delta_{ij}\, a_{i\bullet} - a_{ij}$$

LA is the *Laplacian* of matrix A, and matrix exponentiation (9) can be performed by spectral decomposition of Ψ.

The resulting E is semi-definite positive, with limits $E = \Pi$ for $t \to 0$ (diagonal spatial weights W, expressing complete spatial autarchy), and $E = ff'$ for $t \to \infty$ (constant spatial weights W, expressing complete mobility). Identity $\mathrm{trace}(E(t)) = 1 - t + 0(t^2)$ (Bavaud, 2014) shows t to measure, for $t \ll 1$, the proportion of *distinct* regional pairs in the joint distribution E.

Testing spatial autocorrelation

Under the null hypothesis H_0 of absence of spatial autocorrelation, and under normal approximation, the expected value of the multivariate Moran's I reads (Bavaud, 2014)

$$E_0(I) = \frac{\mathrm{tr}(W) - 1}{n - 1} \qquad \text{where} \quad w_{ij} = \frac{e_{ij}}{f_i}$$

and its the variance reads

$$\mathrm{Var}_0(I) = \frac{2}{n^2 - 1} \left[\mathrm{trace}(W^2) - 1 - \frac{(\mathrm{trace}(W) - 1)^2}{n - 1} \right]$$

Spatial autocorrelation is thus significant at level α if $z := |I - E_0(I)| / \sqrt{\mathrm{Var}_0(I)} \geq u_{1-\frac{\alpha}{2}}$, where $u_{1-\frac{\alpha}{2}}$ is the α^{th} quantile of the standard normal distribution.