

Multilinear Objective Function-based Clustering

Giovanni Rossi

Departments of Computer Science and Engineering - DISI, and Mathematics
University of Bologna, Mura Anteo Zamboni 7, 40126, Bologna, Italy

Keywords: Fuzzy Clustering, Similarity Matrix, Pseudo-Boolean Function, Multilinear Extension, Local Search.

Abstract: The input of most clustering algorithms is a symmetric matrix quantifying similarity within data pairs. Such a matrix is here turned into a quadratic set function measuring cluster score or similarity within data subsets larger than pairs. In general, any set function reasonably assigning a cluster score to data subsets gives rise to an objective function-based clustering problem. When considered in pseudo-Boolean form, cluster score enables to evaluate fuzzy clusters through multilinear extension MLE, while the global score of fuzzy clusterings simply is the sum over constituents fuzzy clusters of their MLE score. This is shown to be no greater than the global score of hard clusterings or partitions of the data set, thereby expanding a known result on extremizers of pseudo-Boolean functions. Yet, a multilinear objective function allows to search for optimality in the interior of the hypercube. The proposed method only requires a fuzzy clustering as initial candidate solution, for the appropriate number of clusters is implicitly extracted from the given data set.

1 INTRODUCTION

Clustering means identifying groups within data, with the intent to have similarity between elements in a same group or cluster, and dissimilarity between elements in different groups. It is important in a variety of fields at the interface of computer science, artificial intelligence and engineering, including pattern recognition and learning, web mining and bioinformatics. In *hard* clustering the sought group structure is a partition of the data set, and clusters are blocks (Aigner, 1997); each data point is in precisely one cluster, with full or unit membership. In *fuzzy* clustering data points distribute $[0, 1]$ -ranged memberships over clusters. This yields more flexibility, which is useful in many applications (Kashef and Kamel, 2010; Valente de Oliveira and Pedrycz, 2007).

In *objective function-based* clustering, the prevailing clusters obtain by maximizing or minimizing an objective function: any solution is mapped into a real quantity, i.e. its efficiency or cost, obtained as the sum over clusters of their own quality. Thus in hard clustering this sum is over blocks, with a *cluster score* set function taking real values over the $2^n - 1$ -set of non-empty data subsets (Roubens, 1982), n being the number of data. How to specify cluster score is the first issue addressed below.

When dealt with in pseudo-Boolean form, cluster score admits a unique polynomial multilinear exten-

sion MLE over n -dimensional unit hypercube $[0, 1]^n$. Such a MLE is a novel and seemingly appropriate measure of the score of fuzzy clusters. This is where to start for the clustering method proposed here. As fuzzy clusterings are collections of fuzzy clusters, attention is placed on those such collections where the data distribute memberships adding up to 1, with optimality found where the sum over fuzzy clusters of their MLE-score is maximal. If global cluster score is evaluated via the MLE of cluster score, then its bounds are on hard clusterings. This expands a known result in pseudo-Boolean optimization (Boros and Hammer, 2002). Clustering is thus approached in terms of set partitioning, with this latter combinatorial optimization problem extended from a discrete to a continuous domain (Pardalos et al., 2006) and solved through a novel local search heuristic.

1.1 Related Work and Approach

Objective function-based fuzzy clustering mainly develops from the fuzzy c -means FcM (Bezdek and Pal, 1992) and the possibilistic c -means PcM (Krishnapuram and Keller, 1996) algorithms. Given n data points in \mathbb{R}^m , with m observed features, both FcM and PcM iteratively act on $c < n$ cluster centers or prototypes, aiming to minimize a cost: the sum over clusters of all distances between cluster elements and their center. For any c centers as input, at each iteration the mem-

berships of data to clusters is re-defined so to minimize the sum of (fuzzy) distances from centers, and next centers themselves are re-calculated so to minimize the sum of distances from (fuzzy) members. In FcM (but not in PcM) membership distributions over the c clusters add up to 1. The iteration stops when two consecutive fuzzy clusterings (specifying c centers and $c \times n$ memberships) are sufficiently close (or coincide). This converges to a local minimum, and given non-convexity of the objective function, the choice of suitable initial cluster centers is crucial. Initial cluster centers have to be arbitrary, and it is harsh to assess whether c is a proper number of clusters for the given data set (Ménard and Eboueya, 2002). Much effort is thus devoted to finding the optimal number of clusters. One approach is to validate fuzzy clusterings obtained at different values of c by means of an index, and then selecting the value of c for the output that scored best on the index (Zahid et al., 2001). This validation may be integrated into the FcM iterations, yielding a validity-guided method (Bensaid et al., 1996). Main *cluster validity indices* are: classification entropy, partition coefficient, uniform data functional, compactness and separation criteria. They may be analyzed comparatively in terms membership distributions over clusters (Pal and Bezdek, 1995). A common idea is that clustering performance is higher the more distributions are concentrated, as this formalizes a non-ambiguous classification (Rezaee et al., 1998; Xie and Beni, 1991).

Recent clustering methods such as *neural gas* (Cottrell et al., 2006), *self organizing maps* (Wu and Chow, 2004), *vector quantization* (Lughofer, 2008; Du, 2010) and *kernel methods* (Shawe-Taylor and Cristianini, 2004) maintain special attention on finding the optimal number of clusters for the given data. In several classification tasks concerning protein structures, text documents, surveys or biological signals, an explicit metric vector space (i.e. \mathbb{R}^m above) is not available. Then, clustering may rely on *spectral methods*, where the $\binom{n}{2}$ similarities within data pairs are the adjacency matrix of a weighted graph, and the (non-zero) eigenvalues and eigenvectors of the associated Laplacian are used for partitioning the data (or vertex) set. Specifically, the sought partition is to be such that lightest edges have endpoints in different blocks, whereas heaviest edges have both endpoints in a same block. Although spectral clustering focuses on hard rather than fuzzy models, still it displays some analogy with the local search method detailed below, as in both cases full membership of data points in prevailing clusters is decided in a single step. In fact, spectral methods mostly focus on some first $c < n$ eigenvalues (in increasing order) (von Luxburg

et al., 2008; Ng et al., 2002), thus constraining the number of clusters. Here, such a constraint is possible as well, although with suitable candidate solution the proposed local search autonomously finds an optimal (unrestricted) number of clusters.

Clustering is here approached by firstly quantifying the cluster score of every non-empty data subset, and secondly in terms of the associated set partitioning combinatorial optimization problem (Korte and Vygen, 2002). Cluster score thus is a set function or, geometrically, a point in \mathbb{R}^{2^n-1} , and rather than measuring a cost (or sum of distances) to be minimized (see above), it measures a worth to be maximized. The idea is to quantify, for every data subset, both internal similarity and dissimilarity with respect to its complement. This resembles the “*collective notion of similarity*” in information-based clustering (Slonim et al., 2005). Objective function-based clustering intrinsically relies on the assumption that every data subset has an associated real-valued worth (or, alternatively, a cost). A main novelty proposed below is to deal with both hard and fuzzy clusters at once by means of the pseudo-Boolean form of set functions. In order to have the same input as in many clustering algorithms, the basic cluster score function provided in the next section obtains from a given similarity matrix, and has polynomial MLE of degree 2 (Boros and Hammer, 2002, pp. 157, 162). This also keeps the computational burden at a seemingly reasonable level.

2 CLUSTER SCORE

Given data set $N = \{1, \dots, n\}$, the input of most clustering algorithms (Xu and Wunsch, 2005) is a symmetric similarity matrix $S \in [0, 1]^{n \times n}$, with S_{ij} quantifying similarity within data pairs $\{i, j\} \subset N$. If data points belong to a Euclidean space, i.e. $N \subset \mathbb{R}^m$, then similarities $S_{ij} = 1 - d(i, j)$ may obtain through a normalized distance $d: N \times N \rightarrow [0, 1]$. Not only pairs but also any data subset $A \subseteq N$ may have a measurable internal similarity (and, possibly, dissimilarity with respect to its complement $A^c = N \setminus A$), interpreted as its score $w(A)$ as a cluster. How to specify set function w_S from given matrix S is addressed hereafter.

For $2^N = \{A: A \subseteq N\}$, collection $\{\zeta(A, \cdot): A \in 2^N\}$ is a linear basis of the vector space \mathbb{R}^{2^N} of real-valued functions w on 2^N , where $\zeta: 2^N \times 2^N \rightarrow \mathbb{R}$ is the element of the incidence algebra (Rota, 1964; Aigner, 1997) of Boolean lattice $(2^N, \cap, \cup)$ defined by $\zeta(A, B) = 1$ if $B \supseteq A$ and $\zeta(A, B) = 0$ if $B \not\supseteq A$. That is, the zeta function. Any w corresponds to linear combination $w(B) = \sum_{A \in 2^N} \mu^w(A) \zeta(A, B) = \sum_{A \subseteq B} \mu^w(A)$ for all $B \in 2^N$, with Möbius inversion $\mu^w: 2^N \rightarrow \mathbb{R}$ given

by (note that \subset denotes strict inclusion) $\mu^w(A) = \sum_{B \subset A} (-1)^{|A \setminus B|} w(B)$ (with $\zeta(B, A) = (-1)^{|A \setminus B|}$), $= w(A) - \sum_{B \subset A} \mu^w(B)$ (recursion, with $w(\emptyset) = 0$). This combinatorial “*analog of the fundamental theorem of the calculus*” (Rota, 1964) yields the unique MLE $f^w : [0, 1]^n \rightarrow \mathbb{R}$ of w , with values $w(B) =$

$$= f^w(\chi_B) = \sum_{A \in 2^N} \left(\prod_{i \in A} \chi_B(i) \right) \mu^w(A) = \sum_{A \subseteq B} \mu^w(A)$$

on vertices, and $f^w(q) = \sum_{A \in 2^N} \left(\prod_{i \in A} q_i \right) \mu^w(A)$ (1)

on any point $q = (q_1, \dots, q_n) \in [0, 1]^n$. Conventionally, $\prod_{i \in \emptyset} q_i := 1$ (Boros and Hammer, 2002, p. 157).

A quadratic MLE is a polynomial of degree 2, i.e. $\mu^w(A) = 0$ if $|A| > 2$. Geometrically, this means that w is a point in a $\binom{n+1}{2}$ -dimensional vector (sub)space, i.e. $w \in \mathbb{R}^{\binom{n+1}{2}}$, as all its $2^n - 1$ values are determined by the values taken by μ^w on the n singletons and on the $\binom{n}{2}$ pairs (with $n + \binom{n}{2} = \binom{n+1}{2}$). Similarity matrix S factually has only $\binom{n}{2}$ valid entries (see below), and trying to exploit them beyond a quadratic form for cluster score w seems clumsy. Also, S is intended precisely to measure such a score S_{ij} for all $\binom{n}{2}$ data pairs. The sought quadratic cluster score function w is thus already defined on such pairs, i.e. $w(\{i, j\}) = S_{ij}$ for all $\{i, j\} \in 2^N$. How to assign scores $w(\{i\})$ to singletons $\{i\}, i \in N$ seems a more delicate matter. If such scores are set equal to the n entries $S_{ii} = 1$ along the main diagonal, then Möbius inversion μ^w takes values $\mu^w(\{i\}) = 1$ on singletons and $\mu^w(\{i, j\}) = S_{ij} - S_{ii} - S_{jj} < 0$ on pairs (while $\mu^w(A) = 0$ for $1 \neq |A| \neq 2$). This is a sufficient (but not necessary) condition for sub-additivity, i.e. $w(A \cup B) - w(A) - w(B) \leq 0$ for all $A, B \in 2^N$ such that $A \cap B = \emptyset$. Then, the trivial partition where each data point is a singleton block is easily checked to be optimal. On the other hand, setting $w(\{i\}) = 0$ for all $i \in N$ yields a Möbius inversion with values $\mu^w(A) \geq 0$ for all $A \in 2^N$, and this is sufficient (but not necessary) for super-additivity, i.e. $w(A \cup B) - w(A) - w(B) \geq 0$ for all $A, B \in 2^N$ with $A \cap B = \emptyset$. Then, it becomes optimal (and again trivial) to place all n data points into a unique grand cluster. An interesting alternative to these two unreasonable situations seems to be

$$w(\{i\}) = \frac{1}{2} - \frac{1}{2(n-1)} \sum_{l \in N \setminus i} S_{il} = \sum_{l \in N \setminus i} \frac{1 - S_{il}}{2(n-1)}.$$

In this way, $1 - S_{il}$ quantifies diversity between $i \in N$ and $l \in N \setminus i$, which must be equally shared among them. The cluster score of singleton $\{i\}$ then is the average of such $n - 1$ diversities $\frac{1 - S_{il}}{2}, l \in N \setminus i$ collected

by i . Möbius inversion takes values $\mu^w(\{i\}) = w(\{i\})$ on singletons and, recursively,

$$\mu^w(\{i, j\}) = \frac{n S_{ij} - 1}{n - 1} - \sum_{l \in N \setminus \{i, j\}} \frac{2 - (S_{il} + S_{jl})}{2(n - 1)}$$

on pairs. Note that the cluster score $w(A)$ of any $A \in 2^N$ does not depend on those $\binom{n-a}{2}$ entries $S_{ll'}$ of matrix S such that $l, l' \in A^c$, where $a = |A|$.

In particular, if $S_{ij} = 1$ for $i \in A, j \in A \setminus i$ and $S_{il} = 0$ for $i \in A, l \in A^c$, then $w(A) = \frac{a(a^2 - 3a + n + 1)}{2(n-1)} =$

$$= a \cdot \frac{n - a}{2(n - 1)} + \binom{a}{2} \cdot \left(1 - \frac{n - a}{n - 1} \right), \text{ with}$$

$w(A) = \frac{1}{2}, 1, \dots, \frac{n^2 - 4n + 5}{2}, \binom{n}{2}$ for $a = 1, 2, \dots, n - 1, n$. In terms of spectral clustering (see above), the corresponding adjacency matrix identifies a subgraph spanned by vertex subset A which is a clique, and where all edges have maximum weight, i.e. 1, with remarkable implications for the eigenvalues of the normalized Laplacian; see (Schaeffer, 2007, p. 42).

This quadratic w , obtained from given similarity matrix S , is conceived as the main example of a cluster score function. Multilinear objective function-based clustering deals with generic set functions, possibly non-quadratic, but reasonably measuring a cluster score of data subsets. In fact, the MLE (1) of set functions w is where to start for the investigation proposed in the remainder of this work. In view of the rich literature on pseudo-Boolean methods (Crama and Hammer, 2011), the MLE of cluster score appears very suitable for evaluating fuzzy clusters, especially in that established definitions of neighborhood and derivative may be expanded to fit the broader setting formalized in the sequel. Specifically, while the n variables of traditional pseudo-Boolean functions range each in $[0, 1]$, the n variables of the novel near-Boolean form (Rossi, 2015) considered here range each in a $2^{n-1} - 1$ -dimensional unit simplex. The reason for this is the purpose to use the MLE of cluster score for evaluating not only fuzzy clusters, but also and most importantly fuzzy clusterings, which are collections $q^1, \dots, q^m \in [0, 1]^n$ of fuzzy clusters, and thus global score is quantifiable as $\sum_{1 \leq k \leq m} f^w(q^k)$. In this respect, note that PcM algorithms allow for membership distributions adding up to quantities < 1 (see above) for handling outliers. These shall be placed each in a singleton block of the partition found here via gradient-based local search. Therefore, membership distributions are like in FcM methods, i.e. ranging in a $2^{n-1} - 1$ -dimensional unit simplex.

3 FUZZY CLUSTERING

A fuzzy clustering is a m -set $q^1, \dots, q^m \in [0, 1]^n$ of fuzzy clusters, with (q_i^1, \dots, q_i^m) being i 's membership distribution. The 2^{n-1} -set of subsets containing each $i \in N$ is $2_i^N = \{A : i \in A \in 2^N\}$, and $\Delta_i = \{(q_i^{A_1}, \dots, q_i^{A_{2^{n-1}}}) \in \mathbb{R}_+^{2^{n-1}} : \sum_{1 \leq k \leq 2^{n-1}} q_i^{A_k} = 1\}$ identifies the corresponding $2^{n-1} - 1$ -dimensional unit simplex, where $\{A_1, \dots, A_{2^{n-1}}\} = 2_i^N$ and $q_i \in \Delta_i$.

Definition 1. A fuzzy cover \mathbf{q} specifies, for each data point $i \in N$, a membership distribution over the 2^{n-1} data subsets $A \in 2_i^N$ containing it. Hence $\mathbf{q} = (q_1, \dots, q_n) \in \Delta_N$, where $\Delta_N = \times_{1 \leq i \leq n} \Delta_i$.

Equivalently, $\mathbf{q} = \{q^A : \emptyset \neq A \in 2^N, q^A \in [0, 1]^n\}$ is a $2^n - 1$ -set whose elements $q^A = (q_1^A, \dots, q_n^A)$ are n -vectors corresponding to non-empty data subsets $\emptyset \neq A \in 2^N$ and specifying a membership q_i^A for each $i \in N$, with $q_i^A \in [0, 1]$ if $i \in A$ while $q_j^A = 0$ if $j \in A^c$. Fuzzy covers thus generalize traditional fuzzy clusterings, as these latter are commonly intended as collections $\{q^1, \dots, q^m\}$ as above where, in addition, for every fuzzy cluster q^k the associated data subset implicitly is $\{i : q_i^k > 0\}$. Conversely, fuzzy covers allow for situations where $0 < |\{i : q_i^A > 0\}| < |A|$ for some $\emptyset \neq A \in 2^N$, although an exactness condition introduced below shows that such cases may be ignored.

Fuzzy covers being collections of points in $[0, 1]^n$, and the MLE f^w of w allowing precisely to evaluate such points, the global score $W(\mathbf{q})$ of any $\mathbf{q} \in \Delta_N$ is the sum over all its elements $q^A, A \in 2^N$ of their own score as quantified by f^w (see (1)). That is,

$$W(\mathbf{q}) = \sum_{A \in 2^N} f^w(q^A) = \sum_{A \in 2^N} \left[\sum_{B \subseteq A} \left(\prod_{i \in B} q_i^A \right) \mu^w(B) \right],$$

$$\text{or } W(\mathbf{q}) = \sum_{A \in 2^N} \left[\sum_{B \supseteq A} \left(\prod_{i \in A} q_i^B \right) \right] \mu^w(A). \quad (2)$$

Example 2. For $N = \{1, 2, 3\}$, consider cluster scores $w(\{1\}) = w(\{2\}) = w(\{3\}) = 0.2$, $w(\{1, 2\}) = 0.8$, $w(\{1, 3\}) = 0.3$, $w(\{2, 3\}) = 0.6$, $w(N) = 0.7$. Membership distributions of data points $i = 1, 2, 3$ over subsets 2_i^N are denoted by $q_1 \in \Delta_1, q_2 \in \Delta_2, q_3 \in \Delta_3$,

$$q_1 = \begin{pmatrix} q_1^1 \\ q_1^{12} \\ q_1^{13} \\ q_1^N \end{pmatrix}, q_2 = \begin{pmatrix} q_2^2 \\ q_2^{12} \\ q_2^{23} \\ q_2^N \end{pmatrix}, q_3 = \begin{pmatrix} q_3^3 \\ q_3^{13} \\ q_3^{23} \\ q_3^N \end{pmatrix}.$$

If $\hat{q}_1^{12} = \hat{q}_2^{12} = 1$, then any membership $q_3 \in \Delta_3$ yields $w(\{1, 2\}) + (q_3^3 + q_3^{13} + q_3^{23} + q_3^N) \mu^w(\{3\}) = w(\{1, 2\}) + w(\{3\}) = 1$. This means that there is a continuum of fuzzy covers achieving maximum score: $W(\hat{q}_1, \hat{q}_2, q_3) = 1$ independently from q_3 . In order to

select the one $\hat{\mathbf{q}} = (\hat{q}_1, \hat{q}_2, \hat{q}_3)$ where $\hat{q}_3^3 = 1$, attention must be placed only on exact ones, defined hereafter.

For any two fuzzy covers $\mathbf{q} = \{q^A : \emptyset \neq A \in 2^N\}$ and $\hat{\mathbf{q}} = \{\hat{q}^A : \emptyset \neq A \in 2^N\}$, define $\hat{\mathbf{q}}$ to be a *shrinking* of \mathbf{q} if there is a subset A , with $\sum_{i \in A} q_i^A > 0$, such that

$$\hat{q}_i^B = \begin{cases} q_i^B & \text{if } B \not\subseteq A \\ 0 & \text{if } B = A \end{cases} \text{ for all } B \in 2^N, i \in N,$$

$$\sum_{B \subset A} \hat{q}_i^B = q_i^A + \sum_{B \subset A} q_i^B \text{ for all } i \in A.$$

In words, a shrinking reallocates the whole membership mass $\sum_{i \in A} q_i^A > 0$ from $A \in 2^N$ to all proper subsets $B \subset A$, involving all and only those data points $i \in A$ with strictly positive membership $q_i^A > 0$.

Definition 3. Fuzzy cover $\mathbf{q} \in \Delta_N$ is exact as long as $W(\mathbf{q}) \neq W(\hat{\mathbf{q}})$ for all shrinkings $\hat{\mathbf{q}}$ of \mathbf{q} .

Proposition 4. If \mathbf{q} is exact, then for all $A \in 2^N$

$$|\{i \in A : q_i^A > 0\}| \in \{0, |A|\}.$$

Proof. For $\emptyset \subset A^+(\mathbf{q}) = \{i : q_i^A > 0\} \subset A \in 2^N$, with $\alpha = |A^+(\mathbf{q})| > 1$, notice that

$$f^w(q^A) = \sum_{B \subseteq A^+(\mathbf{q})} \left(\prod_{i \in B} q_i^A \right) \mu^w(B). \text{ Let shrinking } \hat{\mathbf{q}},$$

with $\hat{q}^{B'} = q^{B'}$ if $B' \not\subseteq 2^{A^+(\mathbf{q})}$, satisfy conditions

$$\sum_{B \in 2_i^N \cap 2^{A^+(\mathbf{q})}} \hat{q}_i^B = q_i^A + \sum_{B \in 2_i^N \cap 2^{A^+(\mathbf{q})}} q_i^B \text{ for all } i \in A^+(\mathbf{q})$$

and $\prod_{i \in B} \hat{q}_i^B = \prod_{i \in B} q_i^B + \prod_{i \in B} q_i^A$ for all $B \in 2^{A^+(\mathbf{q})}, |B| > 1$.

These are $2^\alpha - 1$ equations with $\sum_{1 \leq k \leq \alpha} k \binom{\alpha}{k} > 2^\alpha$ variables $\hat{q}_i^B, B \subseteq A^+(\mathbf{q})$. Thus there is a continuum of solutions, each providing precisely a shrinking $\hat{\mathbf{q}}$

where $\sum_{B \in 2^{A^+(\mathbf{q})}} f^w(\hat{q}^B) = f^w(q^A) + \sum_{B \in 2^{A^+(\mathbf{q})}} f^w(q^B)$.

This entails that \mathbf{q} is not exact. \square

For given w , the global score of any fuzzy cover also attains on many fuzzy clusterings. This justifies the following (in line with standard terminology).

Definition 5. Fuzzy clusterings are exact covers.

The global score of any fuzzy clustering is shown below to also attain on some hard clustering, thereby expanding a result on extremizers of pseudo-Boolean functions (Boros and Hammer, 2002, p. 163).

4 HARD CLUSTERING

Hard clusterings or partitions of N (Aigner, 1997) are fuzzy clusterings where $q_i^A \in \{0, 1\}$ for all $A \in 2^N$ and all $i \in A$. Among the maximizers of any objective function W as above there always exist fuzzy clusterings $(q_1, \dots, q_n) \in \Delta_N$ such that $q_i \in \text{ex}(\Delta_i)$ for all $i \in N$, where $\text{ex}(\Delta_i)$ denotes the 2^{n-1} -set of extreme points of Δ_i . For $\mathbf{q} \in \Delta_N, i \in N$, let $\mathbf{q} = q_i | \mathbf{q}_{-i}$, with $q_i \in \Delta_i$ and $\mathbf{q}_{-i} \in \Delta_{N \setminus i} = \times_{j \in N \setminus i} \Delta_j$. Then, for any w ,

$$\begin{aligned} W(\mathbf{q}) &= \sum_{A \in 2_i^N} f^w(q^A) + \sum_{A' \in 2^{N \setminus 2_i^N}} f^w(q^{A'}) = \\ &= \sum_{A \in 2_i^N} \sum_{B \subseteq A \setminus i} \left(\prod_{j \in B} q_j^A \right) \left(q_i^A \mu^w(B \cup i) + \mu^w(B) \right) + \\ &+ \sum_{A' \in 2^{N \setminus 2_i^N}} \sum_{B' \subseteq A'} \left(\prod_{j \in B'} q_j^{A'} \right) \mu^w(B') \end{aligned}$$

at all $\mathbf{q} \in \Delta_N$ and for all $i \in N$. Now define

$$\begin{aligned} W_i(q_i | \mathbf{q}_{-i}) &= \sum_{A \in 2_i^N} q_i^A \left[\sum_{B \subseteq A \setminus i} \left(\prod_{j \in B} q_j^A \right) \mu^w(B \cup i) \right], \\ W_{-i}(\mathbf{q}_{-i}) &= \sum_{A \in 2_i^N} \left[\sum_{B \subseteq A \setminus i} \left(\prod_{j \in B} q_j^A \right) \mu^w(B) \right] + \\ &+ \sum_{A' \in 2^{N \setminus 2_i^N}} \left[\sum_{B' \subseteq A'} \left(\prod_{j \in B'} q_j^{A'} \right) \mu^w(B') \right], \end{aligned}$$

$$\text{yielding } W(\mathbf{q}) = W_i(q_i | \mathbf{q}_{-i}) + W_{-i}(\mathbf{q}_{-i}). \quad (3)$$

Proposition 6. For all $\mathbf{q} \in \Delta_N$, there are $\underline{\mathbf{q}}, \bar{\mathbf{q}} \in \Delta_N$

$$\text{such that } \begin{cases} (i) W(\underline{\mathbf{q}}) \leq W(\mathbf{q}) \leq W(\bar{\mathbf{q}}) \text{ and,} \\ (ii) \underline{q}_i, \bar{q}_i \in \text{ex}(\Delta_i) \text{ for all } i \in N. \end{cases}$$

Proof. For all $i \in N$ and $\mathbf{q}_{-i} \in \Delta_{N \setminus i}$, define $w_{\mathbf{q}_{-i}} : 2_i^N \rightarrow \mathbb{R}$ by

$$w_{\mathbf{q}_{-i}}(A) = \sum_{B \subseteq A \setminus i} \left(\prod_{j \in B} q_j^A \right) \mu^w(B \cup i). \quad (4)$$

Let $\mathbb{A}_{\mathbf{q}_{-i}}^+ = \arg \max w_{\mathbf{q}_{-i}}$ and $\mathbb{A}_{\mathbf{q}_{-i}}^- = \arg \min w_{\mathbf{q}_{-i}}$, noting that $\mathbb{A}_{\mathbf{q}_{-i}}^+ \neq \emptyset \neq \mathbb{A}_{\mathbf{q}_{-i}}^-$ at all \mathbf{q}_{-i} . Most importantly,

$$W_i(q_i | \mathbf{q}_{-i}) = \sum_{A \in 2_i^N} \left(q_i^A \cdot w_{\mathbf{q}_{-i}}(A) \right) = \langle q_i, w_{\mathbf{q}_{-i}} \rangle, \quad (5)$$

where $\langle \cdot, \cdot \rangle$ denotes scalar product. Thus for given membership distributions of all $j \in N \setminus i$, global score is affected by i 's membership distribution through a scalar product. In order to maximize (or minimize) W by suitably choosing q_i for given \mathbf{q}_{-i} , the whole

of i 's membership mass must be placed over $\mathbb{A}_{\mathbf{q}_{-i}}^+$ (or $\mathbb{A}_{\mathbf{q}_{-i}}^-$), anyhow. Hence there are precisely $|\mathbb{A}_{\mathbf{q}_{-i}}^+| > 0$ (or $|\mathbb{A}_{\mathbf{q}_{-i}}^-| > 0$) available extreme points of Δ_i . The following procedure selects (arbitrarily) one of them.

ROUNDUP(w, \mathbf{q})

Initialize: Set $t = 0$ and $\mathbf{q}(0) = \mathbf{q}$.

Loop: While there is a $i \in N$ with $q_i(t) \notin \text{ex}(\Delta_i)$, set $t = t + 1$ and:

- select some $A^* \in \mathbb{A}_{\mathbf{q}_{-i}(t)}^+$,
- define, for all $j \in N, A \in 2^N$,

$$q_j^A(t) = \begin{cases} q_j^A(t-1) & \text{if } j \neq i \\ 1 & \text{if } j = i \text{ and } A = A^* \\ 0 & \text{otherwise} \end{cases}.$$

Output: Set $\bar{\mathbf{q}} = \mathbf{q}(t)$.

Every change $q_i^A(t-1) \neq q_i^A(t) = 1$ (for any $i \in N, A \in 2_i^N$) induces a non-decreasing variation $W(\mathbf{q}(t)) - W(\mathbf{q}(t-1)) \geq 0$. Hence, the sought $\bar{\mathbf{q}}$ is provided in at most n iterations. Analogously, replacing $\mathbb{A}_{\mathbf{q}_{-i}}^+$ with $\mathbb{A}_{\mathbf{q}_{-i}}^-$ yields the sought minimizer $\underline{\mathbf{q}}$. \square

Remark 7. For $i \in N, A \in 2_i^N$, if all $j \in A \setminus i \neq \emptyset$ satisfy $q_j^A = 1$, then (4) yields $w_{\mathbf{q}_{-i}}(A) = w(A) - w(A \setminus i)$, while $w_{\mathbf{q}_{-i}}(\{i\}) = w(\{i\})$ regardless of \mathbf{q}_{-i} . For quadratic w obtained above from similarity matrix S ,

$$w_{\mathbf{q}_{-i}}(A) = w(\{i\}) + \sum_{j \in A \setminus i} q_j^A \mu^w(\{i, j\}).$$

If the global score of fuzzy clusterings is quantified as the sum over constituents fuzzy clusters of their MLE-score, then for any w there are hard clusterings among both the maximizers and minimizers. This seems crucial because many applications may be modeled in terms of *set partitioning*, and in such a combinatorial optimization problem fuzzy clustering is not feasible. An important example is winner determination in combinatorial auctions (Sandholm, 2002), where a set N of items to be sold must be partitioned into bundles towards revenue maximization. The maximum bid received for each bundle $\emptyset \neq A \subseteq N$ defines the input set function w . The above result entails that if the objective function is multilinearly extended over the continuous domain of fuzzy clusterings, then any found solution can be promptly adapted to the restricted domain of partitions, with no score loss. The problem can thus be approached from a geometrical perspective, allowing for novel search strategies. Partitions $P = \{A_1, \dots, A_{|P|}\} \subset 2^N$ of N are families of pairwise disjoint subsets whose union is N , i.e. $N = \cup_{1 \leq k \leq |P|} A_k$ and $A_k \cap A_l = \emptyset, 1 \leq k < l \leq |P|$. Any P corresponds to the collection $\{\chi_A : A \in P\}$ of

those $|P|$ hypercube vertices identified by the characteristic functions of its blocks (see above). Partitions P can also be seen as $\mathbf{p} \in \Delta_N$ where $p_i^A = 1$ for all $A \in P, i \in A$. The above findings yield the following.

Corollary 8. For any w , some partition P satisfies $W(\mathbf{p}) \geq W(\mathbf{q})$ for all $\mathbf{q} \in \Delta_N$, with $W(\mathbf{p}) = \sum_{A \in P} w(A)$.

Proof. Follows from propositions 4 and 6. \square

A further remark concerns cluster validity (Wang and Zhang, 2007), with focus on those indices that validate fuzzy clusterings by relying exclusively on membership distributions. As already observed, a basic argument is that the more such distributions are concentrated, the less ambiguous is the fuzzy classification. Evidently, hard clusterings provide n distributions each concentrated on a unique extreme point of the associated unit simplex. The above result indicates that if global score is evaluated through MLE, then validation may ignore membership distributions, as the score of any optimal fuzzy clustering also obtains by means of a hard one.

5 LOCAL SEARCH

Defining global maximizers is clearly immediate.

Definition 9. Fuzzy clustering $\hat{\mathbf{q}} \in \Delta_N$ is a global maximizer if $W(\hat{\mathbf{q}}) \geq W(\mathbf{q})$ for all $\mathbf{q} \in \Delta_N$.

Concerning local maximizers, consider a vector $\omega = (\omega_1, \dots, \omega_n) \in \mathbb{R}_{++}^n$ of strictly positive weights, with $\omega_N = \sum_{j \in N} \omega_j$, and focus on the equilibrium (Mas-Colell et al., 1995) of the game where data points are players who strategically choose their memberships distribution $q_i \in \Delta_i$ while being rewarded with fraction $\frac{\omega_i}{\omega_N} W(q_1, \dots, q_n)$ of the global score attained at any strategy profile (q_1, \dots, q_n) .

Definition 10. Fuzzy clustering $\hat{\mathbf{q}} \in \Delta_N$ is a local maximizer if $W_i(\hat{q}_i | \hat{\mathbf{q}}_{-i}) \geq W_i(q_i | \hat{\mathbf{q}}_{-i})$ for all $q_i \in \Delta_i$ and all $i \in N$ (see (3)).

This definition of local maximizer entails that the neighborhood $\mathcal{N}(\mathbf{q}) \subset \Delta_N$ of any $\mathbf{q} \in \Delta_N$ is

$$\mathcal{N}(\mathbf{q}) = \bigcup_{i \in N} \left\{ \tilde{\mathbf{q}} : \tilde{q}_i = \tilde{q}_i | \mathbf{q}_{-i}, \tilde{q}_i \in \Delta_i \right\}.$$

Definition 11. The (i, A) -derivative of W at $\mathbf{q} \in \Delta_N$ is

$$\begin{aligned} \partial W(\mathbf{q}) / \partial q_i^A &= W(\bar{\mathbf{q}}(i, A)) - W(\underline{\mathbf{q}}(i, A)) = \\ &= W_i(\bar{q}_i(i, A) | \bar{\mathbf{q}}_{-i}(i, A)) - W_i(q_i(i, A) | \underline{\mathbf{q}}_{-i}(i, A)), \end{aligned}$$

with $\bar{\mathbf{q}}(i, A) = (\bar{q}_1(i, A), \dots, \bar{q}_n(i, A))$ given by

$$\bar{q}_j^B(i, A) = \begin{cases} q_j^B & \text{for all } j \in N \setminus i, B \in 2_j^N \\ 1 & \text{for } j = i, B = A \\ 0 & \text{for } j = i, B \neq A \end{cases},$$

and $\underline{\mathbf{q}}(i, A) = (\underline{q}_1(i, A), \dots, \underline{q}_n(i, A))$ given by

$$\underline{q}_j^B(i, A) = \begin{cases} q_j^B & \text{for all } j \in N \setminus i, B \in 2_j^N \\ 0 & \text{for } j = i \text{ and all } B \in 2_i^N \end{cases},$$

thus $\nabla W(\mathbf{q}) = \{\partial W(\mathbf{q}) / \partial q_i^A : i \in N, A \in 2_i^N\} \in \mathbb{R}^{n2^{n-1}}$ is the (full) gradient of W at \mathbf{q} . The i -gradient $\nabla_i W(\mathbf{q}) \in \mathbb{R}^{2^{n-1}}$ of W at $\mathbf{q} = q_i | \mathbf{q}_{-i}$ is set function $\nabla_i W(\mathbf{q}) : 2_i^N \rightarrow \mathbb{R}$ defined by $\nabla_i W(\mathbf{q})(A) = w_{\mathbf{q}_{-i}}(A)$ for all $A \in 2_i^N$, where $w_{\mathbf{q}_{-i}}$ is given by (4).

Remark 12. Membership distribution $q_i(i, A)$ is the null one: its 2^{n-1} entries are all 0, hence $q_i(i, A) \notin \Delta_i$.

The setting obtained thus far allows to conceive searching for a local maximizer hard clustering \mathbf{q}^* from given fuzzy clustering \mathbf{q} as initial candidate solution, and while maintaining the whole search within the continuum of fuzzy clusterings. This idea may be specified in alternative ways yielding different local search methods. One possibility is the following.

LOCALSEARCH(w, \mathbf{q})

Initialize: Set $t = 0$ and $\mathbf{q}(0) = \mathbf{q}$, with requirement $|\{i : q_i^A > 0\}| \in \{0, |A|\}$ for all $A \in 2^N$.

Loop 1: While $0 < \sum_{i \in A} q_i^A(t) < |A|$ for a $A \in 2^N$, set $t = t + 1$ and

(a) select a $A^*(t) \in 2^N$ such that

$$\sum_{i \in A^*(t)} w_{\mathbf{q}_{-i}(t-1)}(A^*(t)) \geq \sum_{j \in B} w_{\mathbf{q}_{-j}(t-1)}(B)$$

for all $B \in 2^N$ such that $0 < \sum_{i \in B} q_j^B(t) < |B|$,

(b) for $i \in A^*(t)$ and $A \in 2_i^N$, define

$$q_i^A(t) = \begin{cases} 1 & \text{if } A = A^*(t), \\ 0 & \text{if } A \neq A^*(t), \end{cases}$$

(c) for $j \in N \setminus A^*(t)$ and $A \in 2_j^N$ with $A \cap A^*(t) = \emptyset$, define $q_j^A(t) = q_j^A(t-1) +$

$$\left(w(A) \sum_{\substack{B \in 2_j^N \\ B \cap A^*(t) \neq \emptyset}} q_j^B(t-1) \right) \left(\sum_{\substack{B' \in 2_j^N \\ B' \cap A^*(t) = \emptyset}} w(B') \right)^{-1}$$

(d) for $j \in N \setminus A^*(t)$ and $A \in 2_j^N$ with $A \cap A^*(t) \neq \emptyset$, define

$$q_j^A(t) = 0.$$

Loop 2: While $q_i^A(t) = 1, |A| > 1$ for a $i \in N$ and $w(A) < w(\{i\}) + w(A \setminus i)$, set $t = t + 1$ and define:

$$\begin{aligned} \hat{q}_i^{\hat{A}}(t) &= \begin{cases} 1 & \text{if } |\hat{A}| = 1 \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } \hat{A} \in 2_i^N, \\ q_j^B(t) &= \begin{cases} 1 & \text{if } B = A \setminus i \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } j \in A \setminus i, B \in 2_j^N, \\ q_j^{\hat{B}}(t) &= q_j^{\hat{B}}(t-1) \quad \text{for all } j' \in A^c, \hat{B} \in 2_j^N. \end{aligned}$$

Output: Set $\mathbf{q}^* = \mathbf{q}(t)$.

Both ROUNDUP and LOCALSEARCH yield a sequence $\mathbf{q}(0), \dots, \mathbf{q}(t^*) = \mathbf{q}^*$ where $q_i^* \in \text{ex}(\Delta_i)$ for all $i \in N$. In the former at the end of each iteration t the novel $\mathbf{q}(t) \in \mathcal{N}(\mathbf{q}(t-1))$ is in the neighborhood of its predecessor. In the latter $\mathbf{q}(t) \notin \mathcal{N}(\mathbf{q}(t-1))$ in general, as in $|P| \leq n$ iterations of *Loop 1* a partition $\{A^*(1), \dots, A^*(|P|)\} = P$ is generated. Selected clusters or blocks $A^*(t) \in 2^N$, $t = 1, \dots, |P|$ are any of those where the sum over data points $i \in A^*(t)$ of $(i, A^*(t))$ -derivatives $\partial W(\mathbf{q}(t-1)) / \partial q_i^{A^*(t)}(t-1)$ is maximal. Once a block $A^*(t)$ is selected, then lines (c) and (d) make all data points $j \in N \setminus A^*(t)$ redistribute the entire membership mass currently placed on subsets $A' \in 2_j^N$ with non-empty intersection $A' \cap A^*(t) \neq \emptyset$ over those remaining $A \in 2_j^N$ such that, conversely, $A \cap A^*(t) = \emptyset$. The redistribution is such that each of these latter gets a fraction $w(A) / \sum_{B \in 2_j^N: B \cap A^*(t) = \emptyset} w(B)$ of the newly freed membership mass $\sum_{A' \in 2_j^N: A' \cap A^*(t) \neq \emptyset} q_j^{A'}(t-1)$. The subsequent *Loop 2* checks whether the partition generated by *Loop 1* may be improved by extracting some outliers from existing blocks and putting them in singleton blocks of the final output. An outlier basically is a data point displaying very unusual features. In the limit, cluster score w may be such that for some data points $i \in N$ global score decreases when i joins any cluster $A \in 2_i^N$, $|A| > 1$, that is to say $w(A) - w(A \setminus i) - w(\{i\}) = \sum_{B \in 2^A \setminus 2^A \setminus i: |B| > 1} \mu^w(B) < 0$.

Proposition 13. *Output \mathbf{q}^* of LOCALSEARCH(W, \mathbf{q}) is a local maximizer.*

Proof. It is plain that the output corresponds to a partition P . With the notation of corollary 8 in section 4, $\mathbf{q}^* = \mathbf{p}$. Accordingly, any data point $i \in N$ is either in a singleton cluster $\{i\} \in P$ or else in a cluster $A \in P, i \in A$ such that $|A| > 1$. In the former case, any membership reallocation deviating from $p_i^{\{i\}} = 1$, given memberships $p_j, j \in N \setminus i$, yields a cover (fuzzy or hard) where global score is the same as at \mathbf{p} , because $\prod_{j \in B \setminus i} p_j^B = 0$ for all $B \in 2_i^N \setminus A$ (see example 2 above). In the latter case, any membership reallocation q_i deviating from $p_i^A = 1$ (given memberships $p_j, j \in N \setminus i$) yields a cover which is best seen by distinguishing between $2_i^N \setminus A$ and A . Also recall that $w(A) - w(A \setminus i) = \sum_{B \in 2^A \setminus 2^A \setminus i} \mu^w(B)$. Again, all membership mass $\sum_{B \in 2_i^N \setminus A} q_i^B > 0$ simply collapses on singleton $\{i\}$ because $\prod_{j \in B \setminus i} p_j^B = 0$ for all $B \in 2_i^N \setminus A$. Therefore, $W(\mathbf{p}) - W(q_i | \mathbf{p}_{-i}) = w(A) - w(\{i\}) +$

$$- \left(q_i^A \sum_{B \in 2^A \setminus 2^A \setminus i: |B| > 1} \mu^w(B) + \sum_{B' \in 2^A \setminus i} \mu^w(B') \right) =$$

$$= (p_i^A - q_i^A) \sum_{B \in 2^A \setminus 2^A \setminus i: |B| > 1} \mu^w(B).$$

Now assume that \mathbf{q} is *not* a local maximizer, i.e. $W(\mathbf{p}) - W(q_i | \mathbf{p}_{-i}) < 0$. Since $p_i^A - q_i^A > 0$ (because $p_i^A = 1$ and $q_i \in \Delta_i$ is a deviation from p_i), then

$$\sum_{B \in 2^A \setminus 2^A \setminus i: |B| > 1} \mu^w(B) = w(A) - w(A \setminus i) - w(\{i\}) < 0.$$

Hence \mathbf{q} cannot be the output of *Second Loop*. \square

In local search methods, the chosen initial candidate solution determines what neighborhoods shall be visited. The range of the objective function in a neighborhood is a set of real values. In a neighborhood $\mathcal{N}(\mathbf{p})$ of a hard clustering \mathbf{p} or partition P only those $\sum_{A \in P: |A| > 1} |A|$ data points $i \in A$ in non-singleton blocks $A \in P, |A| > 1$ can modify global score by reallocating their membership. In view of the above proof, the only admissible variations obtain by deviating from $p_i^A = 1$ with an alternative membership distribution q_i such that $q_i^A \in [0, 1)$, with $W(q_i | \mathbf{p}_{-i}) - W(\mathbf{p})$ equal to $(q_i^A - 1) \sum_{B \in 2^A \setminus 2^A \setminus i: |B| > 1} \mu^w(B) + (1 - q_i^A) w(\{i\})$. Hence, choosing partitions as initial candidate solutions of LOCALSEARCH is evidently poor. A sensible choice should conversely allow the search to explore different neighborhoods where the objective function may range widely. A simplest example of such an initial candidate solution is uniform distribution $q_i^A = 2^{1-n}$. On the other hand, the input of local search fuzzy clustering algorithms is commonly desired to be close to a global optimum, i.e. a maximizer in the present setting. This translates here into the idea of defining a suitable input by means of cluster score function w . Along this line, consider $q_i^A = w(A) / \sum_{B \in 2_i^N} w(B)$, yielding $\frac{q_i^A}{q_i^B} = \frac{w(A)}{w(B)} = \frac{q_j^A}{q_j^B}$ for all $A, B \in 2_i^N \cap 2_j^N$ and all $i, j \in N$.

With a suitable initial candidate solution, the search may be restricted to explore only a maximum number of fuzzy clusterings, thereby containing (together with the quadratic MLE of cluster score w) the computational burden. In particular, if $\mathbf{q}(0)$ is the finest partition $\{\{1\}, \dots, \{n\}\}$ or $q_i^{\{i\}}(0) = 1$ for all $i \in N$, then the search does not explore any neighborhood at all, and such an input coincides with the output. More reasonably, let $\mathbb{A}_{\mathbf{q}}^{\max} = \{A_1, \dots, A_k\}$ denote the collection of maximal data subsets where input memberships are strictly positive. That is, $q_i^{A_{k'}} > 0$ for all $i \in A_{k'}, 1 \leq k' \leq k$ as well as $q_j^B = 0$ for all $B \in 2^N \setminus (2^{A_1} \cup \dots \cup 2^{A_k})$ and all $j \in B$. Then, the output shall be a partition P each of whose blocks $A \in P$ satisfies $A \subseteq A_{k'}$ for some $1 \leq k' \leq k$. Hence, by suitably choosing the input \mathbf{q} , LOCALSEARCH outputs

a partition with no less than some maximum desired number $k(\mathbf{q})$ blocks.

6 CONCLUSIONS

This paper approaches objective function-based fuzzy clustering by firstly eliciting a real-valued cluster score function, quantifying the positive worth of data subsets in the given classification problem. Clustering is next interpreted in terms of combinatorial optimization via set partitioning. The proposed gradient-based local search relies on a novel expansion of the MLE of near-Boolean functions (Rossi, 2015) over the product of n simplices, each of which is $2^{n-1} - 1$ -dimensional, n being the number of data. The method needs not the input to specify a desired number of clusters, as this latter is determined autonomously through optimization, and applies to any classification problem, handling data sets not necessarily included in a Euclidean space: proximities between data points and within clusters may be quantified in any conceivable way, including information theoretic measurement (Pirr6 and Euzenat, 2010).

REFERENCES

- Aigner, M. (1997). *Combinatorial Theory*. Springer. Reprint of the 1979 Edition.
- Bensaid, A., Hall, L., Bezdek, J., Clarke, L., Silbiger, M., Arrington, J., and Murtagh, R. (1996). Validity-guided (re)clustering with applications to image segmentation. *IEEE Trans. on Fuzzy Sys.*, 4(2):112–123.
- Bezdek, J. and Pal, S. (1992). *Fuzzy Models for Pattern Recognition*. IEEE Press.
- Boros, E. and Hammer, P. (2002). Pseudo-Boolean optimization. *Discrete App. Math.*, 123:155–225.
- Cottrell, M., Hammer, B., Hasenfuß, A., and Villmann, T. (2006). Batch and median neural gas. *Neural Networks*, 19(6-7):762–771.
- Crama, Y. and Hammer, P. L. (2011). *Boolean Functions: Theory, Algorithms, and Applications*. Cambridge University Press.
- Du, K.-L. (2010). Clustering: a neural network approach. *Neural Networks*, 23:89–107.
- Kashef, R. and Kamel, M. S. (2010). Cooperative clustering. *Pattern Recognition*, 43(6):2315–2329.
- Korte, B. and Vygen, J. (2002). *Combinatorial Optimization. Theory and Algorithms*. Springer.
- Krishnapuram, R. and Keller, J. (1996). The possibilistic c-means algorithm: insights and recommendations. *IEEE Transactions on Fuzzy Systems*, 4(3):148–158.
- Lughofer, E. (2008). Extensions of vector quantization for incremental clustering. *Pattern Recognition*, 41:995–1011.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- Ménard, M. and Eboueya, M. (2002). Extreme physical information and objective function in fuzzy clustering. *Fuzzy Sets and Systems*, 128:285–303.
- Ng, A. Y., Jordan, M. I., and Weiss, Y. (2002). On spectral clustering: analysis and an algorithm. In Dietterich, T. G., Becker, S., and Ghahramani, Z., editors, *Advances in Neural Information Processing Systems 14*, volume 2, pages 849–856. MIT Press.
- Pal, N. and Bezdek, J. (1995). On cluster validity for the fuzzy c-means model. *IEEE Transactions on Fuzzy Systems*, 3(3):370–379.
- Pardalos, P., Prokopyev, O., and Busygin, S. (2006). Continuous approaches for solving discrete optimization problems. In Appa, G., Pitsoulis, L., and Williams, H., editors, *Handbook on Modeling for Discrete Optimization*, pages 39–60. Springer.
- Pirr6, G. and Euzenat, J. (2010). A feature and information theoretic framework for semantic similarity and relatedness. In *Proceedings of The Semantic Web Conference ISWC 2010*, pages 615–630. LNCS 6496.
- Rezaee, M., Lelieveldt, B., and Reiber, J. (1998). A new cluster validity index for the fuzzy c-means. *Pattern Recognition Letters*, 19:237–246.
- Rossi, G. (2015). *Continuous set packing problems and near-Boolean functions*. arXiv 1509.07986v1. Submitted to ICPRAM 2016.
- Rota, G.-C. (1964). On the foundations of combinatorial theory I: theory of Möbius functions. *Z. Wahrscheinlichkeitsrechnung u. verw. Geb.*, 2:340–368.
- Roubens, M. (1982). Fuzzy clustering algorithms and their cluster validity. *European Journal of Operational Research*, 10(3):294–301.
- Sandholm, T. (2002). Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, (135):1–54.
- Schaeffer, S. E. (2007). Graph clustering. *Computer Science Review*, 1(1):27–64.
- Shawe-Taylor, J. and Cristianini, N. (2004). *Kernel Method for Pattern Analysis*. Cambridge University Press.
- Slonim, N., Atwal, S. G., Tkačik, G., and Bialek, W. (2005). Information-based clustering. *PNAS*, 102(51):18297–18302.
- Valente de Oliveira, J. and Pedrycz, W. (2007). *Advances in fuzzy clustering and its applications*. Wiley.
- von Luxburg, U., Belkin, M., and Bousquet, O. (2008). Consistency of spectral clustering. *The Annals of Statistics*, 36(2):555–586.
- Wang, W. and Zhang, Y. (2007). On fuzzy cluster validity indices. *Fuzzy Sets and Systems*, 158:2095–2117.
- Wu, S. and Chow, T. W. S. (2004). Clustering of the self-organizing map using a cluster validity index based on inter-cluster and intra-cluster density. *Pattern Recognition*, 37:175–188.
- Xie, X. and Beni, G. (1991). Validity measure for fuzzy clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):841–847.

- Xu, R. and Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Trans. on Neural Net.*, 16(3):645–678.
- Zahid, N., Abouelala, O., Limouri, M., and Essaid, A. (2001). Fuzzy clustering based on k-nearest-neighbours rule. *Fuzzy Sets and Sys.*, 120:239–247.