



Regularized bi-directional co-clustering

Séverine Affeldt¹ · Lazhar Labiod¹ · Mohamed Nadif¹

Received: 18 June 2020 / Accepted: 1 March 2021 / Published online: 10 April 2021
© Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

The simultaneous clustering of documents and words, known as *co-clustering*, has proved to be more effective than one-sided clustering in dealing with sparse high-dimensional datasets. By their nature, text data are also generally unbalanced and directional. Recently, the von Mises–Fisher (vMF) mixture model was proposed to handle unbalanced data while harnessing the directional nature of text. In this paper, we propose a general co-clustering framework based on a matrix formulation of vMF model-based co-clustering. This formulation leads to a flexible framework for text co-clustering that can easily incorporate both *word–word* semantic relationships and *document–document* similarities. By contrast with existing methods, which generally use an additive incorporation of similarities, we propose a *bi-directional multiplicative* regularization that better encapsulates the underlying text data structure. Extensive evaluations on various real-world text datasets demonstrate the superior performance of our proposed approach over baseline and competitive methods, both in terms of clustering results and co-cluster topic coherence.

Keywords Co-clustering · Regularization · Information retrieval · Text mining

1 Introduction

The simultaneous partitioning of features and objects into consistent homogeneous blocks, referred as to *co-clusters*¹, is a successful extension of one-sided clustering that can make large datasets easier to analyze (Hartigan 1972; Bock 1979; Govaert 1983; Vichi 2001; Van Mechelen et al. 2004; Rocci and Vichi 2008; Govaert and Nadif 2008, 2013; Bock 2020). Starting from a data matrix, a co-cluster can be defined as a submatrix whose elements have a particular pattern in common. The basic idea behind co-clustering is to identify a structure that is shared by objects and features through their permutations. A variety of co-clustering methods have been applied in different areas, such as in bioinformatics

(Madeira and Oliveira 2004; Cho and Dhillon 2008; Tanay et al. 2005; Hanczar and Nadif 2012) to group genes and experimental conditions, in collaborative filtering (Hofmann and Puzicha 1999; Deodhar and Ghosh 2010) to group users and items, and in text mining (Dhillon 2003; Ailem et al. 2017a; Govaert and Nadif 2018; Salah and Nadif 2019; Role et al. 2019) to group words² and documents. Through its ability to relate rows and columns, co-clustering generally gives better results than clustering along a single dimension. Besides, co-clustering makes an implicit adaptive dimensionality reduction that allows the use of efficient scalable algorithms for high-dimensional sparse text data. This is crucial in text mining, since the exponential growth of online documents has created an urgent need for effective methods in handling and interpreting high-dimensional sparse *document-term* matrices, i.e., matrices where documents are represented in the space of terms, and vice versa. Most importantly, text co-clustering can identify the most discriminating words that characterize topics in document classes.

¹ Given a data matrix $\mathbf{X} = (x_{ij}), i \in I, j \in J$, a co-cluster is a submatrix defined by $I_k \times J_\ell (I_k \subseteq I, J_\ell \subseteq J)$.

✉ Séverine Affeldt
severine.affeldt@u-paris.fr

Lazhar Labiod
lazhar.labiod@u-paris.fr

Mohamed Nadif
mohamed.nadif@u-paris.fr

¹ Université de Paris, CNRS, Centre Borelli, Paris F-75006, France

² The generally understood difference between *words* and *terms* is that *terms* are *words* used in a particular specialized field. The *words* that we are concerned with in co-clustering can in most cases also be qualified as *terms*, and consequently, we use *words* and *terms* interchangeably in this paper.

Standard text-focused co-clustering approaches seek to relate documents and words. They do not usually attempt to incorporate side information such as semantic relationships between words, or similarities in document content. The clustering of documents relating to the same topics might nevertheless benefit from additional information about word similarities, since these documents can be expected to contain semantically related terms. Conversely, word clustering might usefully harness side information about (similarities in) document content, given that it is the documents that provide the context for the words. Side information on document latent space *and* on word latent space could together improve the co-clustering of document-text data.

2 Related work

Inspired by the recent success of neural word embedding models, Ailem et al. (2017b) proposed performing NMF (Non-negative Matrix Factorization) jointly on the document-word and word-context matrices, with shared word factors. Recently, an extension of NMTF-based (Non-negative Matrix Tri-Factorization) co-clustering, namely WC-NMTF (Word Co-Occurrence regularized NMTF) (Salah et al. 2018), a technique that takes account of semantic relationships between terms, has successfully been applied on various text datasets. As well as being high-dimensional and sparse, text data are also heavily unbalanced, and co-clustering methods that focus on document-term matrices need to take this into account. The DCC algorithm (Directional Co-clustering with a Conscience) (Salah and Nadif 2019) has been shown to be particularly suited to tackling this issue. DCC uses the von Mises–Fisher (vMF) mixture model and introduces a *conscience* mechanism (DeSieno 1988; Ahalt et al. 1990) to avoid empty or highly unbalanced clusters (Banerjee and Ghosh 2004). It exploits the fact that text data are naturally *directional*, which means that only the directions of data vectors are relevant, and not their magnitude (Mardia and Jupp 2009). In contrast to WC-NMTF, DCC does not use any regularization.

In this work, we harness the directional property of text data and describe a *Regularized Bi-Directional Co-clustering* (RBDCO) algorithm for document-term data. The bi-directional aspect of our approach resides in the use of side information for the two dimensions of the document-term matrix. The primary contribution of this work is a general framework based on a matrix formulation of vMF-based co-clustering. A significant outcome of this novel formulation is a very rich, flexible framework for text co-clustering that allows an easy multiplicative regularization on both the *word–word* semantic relationships and the *document–document* content similarities. In contrast to existing methods, which generally rely on additive incorporation

of similarities, we propose a *bi-directional multiplicative* regularization that better encapsulates the underlying text data structure. Another contribution of this work is an original method for evaluating the coherence of word clusters. Experimental results on various real-life datasets provide clear empirical evidence of the effectiveness of our co-clustering framework.

The rest of the paper is organized as follows: After reviewing the von Mises–Fisher-based clustering method in Sect. 3, we introduce a matrix view of a derived co-clustering algorithm, namely Directional Co-Clustering with a Conscience (DCC), in Sect. 4. We then show in Sect. 5 how a generalized regularization framework can be built from the von Mises–Fisher model while taking into account the directional property of text data, and this section also looks at how our generalized framework is linked to a variety of other co-clustering approaches. Section 6 is devoted to comparative numerical experiments that demonstrate the effectiveness of our generalized regularization framework. We conclude and suggest future paths in Sect. 7.

2.1 Notation

Let $\mathbf{X} = (x_{ij})$ be a data matrix of size $n \times d$, $x_{ij} \in \mathbb{R}$. The i th row of this matrix is represented by a vector $\mathbf{x}_i = (x_{i1}, \dots, x_{id})^\top$, where \top denotes the transpose. The partition of the set of rows into g clusters can be represented by a classification matrix \mathbf{Z} of elements z_{ik} in $\{0, 1\}$ satisfying $\sum_{k=1}^g z_{ik} = 1$. We denote by $z_{\cdot k} = \sum_i z_{ik}$ the cardinality of the k th row cluster. The notation $\mathbf{z} = (z_1, \dots, z_n)^\top$, where $z_i \in \{1, \dots, g\}$ corresponds to the cluster label of i , will be also used. Similarly, the notations $\mathbf{W} = (w_{jk})$, $w_{jk} \in \{0, 1\}$ satisfying $\sum_{k=1}^g w_{jk} = 1$, $\mathbf{w} = (w_1, \dots, w_d)$, where $w_j \in \{1, \dots, g\}$ represents the cluster label of j represented by the vector \mathbf{x}^j , and $w_{\cdot k} = \sum_j w_{jk}$ the cardinality of the k th column cluster, will be used to represent the partition of the set of columns.

3 Directional co-clustering

Mixture models have undoubtedly made a very useful contribution to clustering in that they offer considerable flexibility (McLachlan and Peel 2004). A mixture of von Mises–Fisher (vMF) distributions can be a wise choice (Banerjee et al. 2005; Salah and Nadif 2017b) when dealing with directional data distributed on a unit hypersphere \mathbb{S} . In fact, this model is one of the most appropriate models for clustering high-dimensional sparse data such as the document-term matrices encountered in text mining. In this kind of application, it has been empirically demonstrated that vMF-based clustering methods perform better than a number of existing




Fig. 1 Graphical model. The parameters μ_{z_i} and κ_{z_i} are the mean direction and concentration parameter of vMF distribution $f(\mathbf{x}_i | \mu_{z_i}^w, \kappa_{z_i}) = c_d(\kappa) \exp[\kappa_{z_i} (\mu_{z_i}^w)^\top \mathbf{x}_i]$, respectively. The normalization term takes the following form $c_d(\kappa) = \kappa^{\frac{d}{2}-1} (2\pi)^{\frac{d}{2}} I_{\frac{d}{2}-1}(\kappa)$, where $I_r(\kappa)$ represents the modified Bessel function of the first kind and order r . For more details on the vMF distribution, the reader can refer to Mardia and Jupp (2009)

approaches, see, e.g., (Zhong and Ghosh 2005; Gopal and Yang 2014).

In Salah and Nadif (2017a, 2019), the authors proposed a vMF mixture model for co-clustering. The graphical model is depicted in Fig. 1, and its probability density function is given by

$$f(\mathbf{x}_i | \Theta) = \sum_{k=1}^g \pi_k f_k(\mathbf{x}_i | \mu_{z_i}^w, \kappa_{z_i}),$$

where $\Theta = \{\pi = (\pi_1, \dots, \pi_g), \mu^w = (\mu_1^w, \dots, \mu_g^w), \kappa_1, \dots, \kappa_g\}$. Note that μ^w depends on w , i.e., $w_j = k$ if the j th column belongs to k th cluster, that is ‘‘associated’’ with the k th row cluster.

Note that with this model the d -dimensional centroids $\mu_k^w = (\mu_{k1}, \dots, \mu_{k1}, \dots, \mu_{kg}, \dots, \mu_{kg})^\top$ such that μ_{kh} is repeated w_h times, and $\mu_{kh} = 0$ for all $k \neq h$ are assumed to be orthonormal. The parameter κ_k denotes the concentration of the k th distribution. The proportion of points \mathbf{x}_i generated from the k th component is denoted by the parameter π_k , such that $\sum_k \pi_k = 1$ and $\pi_k > 0, \forall k \in \{1, \dots, g\}$. The complete data log-likelihood is thereby given by

$$L_c(\Theta | \mathbf{X}, \mathbf{Z}) = \sum_k z_{.k} \log \pi_k + \sum_k z_{.k} \log(c_d(\kappa_k)) + \sum_{i,k} z_{ik} \kappa_k (\mu_k^w)^\top \mathbf{x}_i.$$

Assuming that all the mixing proportions are equal, i.e., $\pi_k = \frac{1}{g}, \forall k$ (this does not penalize the quality of clustering as a result of the conscience mechanism) and for high dimensionality, i.e., large order $r = d/2 - 1$, a small κ_k (due to the sparsity) gives $4(r + 1) + \kappa_k^2 \approx 4(r + 1)$ and then $\log c_d(\kappa_k) \approx -\frac{d}{2} \log 2\pi - \log c$ where $c = \frac{4(r+1)}{2^{r+2}(r+1)!}$; for details, the reader can refer to Salah and Nadif (2017a). Thus, $L_c(\Theta | \mathbf{X}, \mathbf{Z})$ becomes

$$L_c(\Theta | \mathbf{X}, \mathbf{Z}) = \sum_{i,k} z_{ik} \kappa_k (\mu_k^w)^\top \mathbf{x}_i + \text{constant}. \tag{1}$$

The concentration parameter κ_k is made inversely proportional to the root square of the number of elements in cluster k , i.e., $\kappa_k = 1/\sqrt{z_{.k}}$ where the row assignments are done by maximizing a weighted Skmeans-like criterion where the weights $1/\sqrt{z_{.k}}$ ($k \in \{1, \dots, g\}$) discourage the absorption of new objects by larger clusters. This is also the case for the column assignments, where $w_{.k}$ is the cardinality of the k th column cluster and $\mu_k^z = (\mu_{k1}, \dots, \mu_{k1}, \dots, \mu_{kg}, \dots, \mu_{kg})^\top$ its n -dimensional centroid. To sum up, following Salah and Nadif (2017a) we have,

$$L_c(\Theta | \mathbf{X}, \mathbf{Z}) \equiv \sum_{i,k} z_{ik} \frac{1}{\sqrt{z_{.k}}} (\mu_k^w)^\top \mathbf{x}_i \equiv \sum_{j,k} w_{jk} \frac{1}{\sqrt{w_{.k}}} (\mu_k^z)^\top \mathbf{x}^j. \tag{2}$$

$A \equiv B$ means that optimizing A is equivalent to optimizing B . The maximizing of mean directions (2) is defined as follows: $\mu_{kh} = 1/\sqrt{w_{.k}}$ if $k = h$, and $\mu_{kh} = 0$ for all $k \neq h$. Similarly, we deduce μ_k^w from $1/\sqrt{z_{.k}}$. Note that Θ is now reduced to μ^w and μ^z the centers of row and column clusters.

The authors have derived a co-clustering algorithm that we refer to as *Directional Co-clustering with a Conscience* (DCC), tailored to high-dimensional sparse data (Alg. 1). The DCC algorithm intertwines row and column clusterings at each step so as to optimize $L_c(\Theta | \mathbf{X}, \mathbf{Z})$. Integrating the conscience mechanism makes it possible to avoid highly skewed solutions with empty or very small/large clusters. Applied on unbalanced document-term matrices, DCC proves more suitable than most existing co-clustering approaches for handling directional data distributed on the surface of a unit-hypersphere.

4 Matrix view of the DCC model

In this section, we propose a matrix formulation of DCC. To this end, we first make use of the matrix formulation of $\mu^w = (\mu_1^w, \dots, \mu_k^w, \dots, \mu_g^w) \in \mathbb{R}^{d \times g}$ and $\mu^z = (\mu_1^z, \dots, \mu_k^z, \dots, \mu_g^z) \in \mathbb{R}^{n \times g}$. Let us consider the binary classification matrices $\mathbf{Z} \in \{0, 1\}^{n \times g}$ and $\mathbf{W} \in \{0, 1\}^{d \times g}$, where the cluster sizes of \mathbf{Z} and \mathbf{W} are on the diagonal of $\mathbf{D}_z = \mathbf{Z}^\top \mathbf{Z}$ and $\mathbf{D}_w = \mathbf{W}^\top \mathbf{W}$, respectively. We therefore have

$$\mu^w = \mathbf{W} \mathbf{D}_w^{-0.5} = \tilde{\mathbf{W}}. \quad \text{and} \quad \mu^z = \mathbf{Z} \mathbf{D}_z^{-0.5} = \tilde{\mathbf{Z}}. \tag{3}$$

Using the above matrix formulations, and given a document-term matrix \mathbf{X} , the optimization of the complete data log-likelihood of \mathbf{X} $L_c(\Theta | \mathbf{X}, \mathbf{Z})$ in (2) leads to

Algorithm 1 Directional Co-clustering with a Conscience (DCC).

Input: \mathbf{X} ($x_i \in \mathbb{S}^{d-1}$ the unit hypersphere), g the number of co-clusters.
Output: \mathbf{Z} , \mathbf{W} and Θ
Steps:
 Initialization: $\Theta \leftarrow \Theta^{(0)}$;
repeat
 1. Assignment of objects:
 for $i = 1$ **to** n **do**
 $z_i = \arg \max_{k'} \frac{1}{\sqrt{z_{i,k'}}} \cos(\boldsymbol{\mu}_{k'}^w, \mathbf{x}_i)$
 end for
 2. Assignment of features:
 for $j = 1$ **to** d **do**
 $w_j = \arg \max_{k'} \frac{1}{\sqrt{w_{j,k'}}} \cos(\boldsymbol{\mu}_{k'}^z, \mathbf{x}^j)$
 end for
 3. Computation of μ_{kk} 's maximizing (2):
 for $k = 1$ **to** g **do**
 $\mu_{kk} \leftarrow \frac{1}{\sqrt{w_{j,k}}}$
 end for
until the (2) value change is small or there is no change

$$\sum_{i,k} z_{ik} \frac{1}{\sqrt{z_{i,k}}} (\boldsymbol{\mu}_{k'}^w)^\top \mathbf{x}_i \equiv Tr(\mathbf{Z}^\top \mathbf{X} \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5}) \tag{4}$$

$$\equiv Tr(\mathbf{W}^\top \mathbf{X}^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5}).$$

In virtue of (3), the formulas for updating the algorithm DCC can be rewritten in the matrix form:

$$\mathbf{Z} = \mathbf{Binarize}(\mathbf{X} \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5}) \tag{5}$$

and

$$\mathbf{W} = \mathbf{Binarize}(\mathbf{X}^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5}),$$

where $\mathbf{Binarize}(\mathbf{B})$, means $\forall i; \mathbf{b}_{ik} = \arg \max_{k'} \mathbf{b}_{ik'}$. The update rules show the mutual interaction between the set of documents and the set of words. If a word w is common to many documents associated with a co-cluster C_i , then the word w will be associated with the co-cluster C_i . Conversely, if a document contains many words that are associated with a co-cluster C_i , then the document will be associated with the co-cluster C_i . To find the desired solution for the partitions \mathbf{Z} and \mathbf{W} , we can alternate the two rules (5) until a fixed point is reached.

5 Regularized bi-directional co-clustering

Text data co-clustering relies on the duality between the document and word spaces, i.e., documents can be grouped based on their distribution with respect to words, while words can be grouped based on their distribution with respect to documents. Existing co-clustering algorithms generally rely on the input *document-term* matrix \mathbf{X} . While some of them con-

sider also pure *word-word* semantic correlations, such as proposed by Salah et al. (2018), co-clustering methods fail to consider side information arising from both *word-word* semantic correlations and *document-document* similarities. To fill this gap, we propose a *Regularized Bi-directional Co-clustering* (RBDCo) based on an appropriate matrix formulation. We construct two similarity matrices—the first, \mathbf{S}_r , for similarities in document content, and the second, \mathbf{S}_c , for semantic correlations between words: see Sect. 5.5—in order to exploit hidden structures in documents and words. Our co-clustering method is then formulated as an iterative matrix multiplication process with two similarity matrices as regularizers, which means that the partitions of documents and words need to be smoothed with respect to document similarities and semantic correlations of words.

Formally, let us consider the block matrix $[\mathbf{Z} \ \mathbf{W}]^\top$. Utilizing the diagonal structure of $\begin{bmatrix} 0 & \mathbf{X} \\ \mathbf{X}^\top & 0 \end{bmatrix}$, we can write the update rules of the aforementioned DCC as an iterative matrix multiplication procedure based on the appropriate block matrices:

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{W} \end{bmatrix} \leftarrow \begin{bmatrix} 0 & \mathbf{X} \\ \mathbf{X}^\top & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \\ \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \end{bmatrix} = \begin{bmatrix} \mathbf{X} \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \\ \mathbf{X}^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \end{bmatrix}. \tag{6}$$

This formulation³ clearly shows how DCC utilizes the duality between document and word spaces. The document clustering \mathbf{Z} is derived as a weighted projection of the data matrix \mathbf{X} on the subspace spanned by the word partition \mathbf{W} . Similarly, the word partition is derived as a weighted projection of the data matrix \mathbf{X} on the subspace spanned by the document partition \mathbf{Z} .

5.1 RBDCo method

We propose a regularized bi-directional data co-clustering method, RBDCo, that draws advantage from our block matrix formulation of DCC (Eq. 6) and harnesses two regularized data matrices, \mathbf{M}_z and \mathbf{M}_w , with values in $\{\mathbf{X}, \mathbf{S}_r \mathbf{X}, \mathbf{X} \mathbf{S}_c, \mathbf{S}_r \mathbf{X} \mathbf{S}_c\}$ (see Sect. 5.5). The objective of RBDCo is to optimize the following trace criterion:

$$J_{RBDCo} \equiv \frac{1}{2} Tr \left(\begin{bmatrix} \mathbf{Z} \\ \mathbf{W} \end{bmatrix}^\top \begin{bmatrix} 0 & \mathbf{M}_z \\ \mathbf{M}_w^\top & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \\ \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \end{bmatrix} \right)$$

$$\equiv \frac{1}{2} Tr \left(\mathbf{Z}^\top \mathbf{M}_z \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \right) + Tr(\mathbf{W}^\top \mathbf{M}_w^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5})$$

$$\equiv \frac{1}{2} Tr \left(\tilde{\mathbf{Z}}^\top (\mathbf{M}_z + \mathbf{M}_w) \tilde{\mathbf{W}} \right). \tag{7}$$

³ to simplify notation, in the rest of the paper the symbol \leftarrow in the updating rules for \mathbf{Z} and \mathbf{W} will indicate that the function $\mathbf{Binarize}(\cdot)$ is applied to the formulas of both \mathbf{Z} and \mathbf{W} .

Algorithm 2 Regularized Bi-Directional Co-Clustering (RBDCo).

Input: \mathbf{X} ($\mathbf{x}_i \in \mathbb{S}^{d-1}$), g number of co-clusters, $\mathbf{S}_r, \mathbf{S}_c$
Output: partitions \mathbf{Z} and \mathbf{W}
Initialization: random initialization of \mathbf{Z} and \mathbf{W}
repeat
 1. Assignment of objects (8)
 • $\mathbf{Z} \leftarrow \mathbf{M}_z \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5}$
 • **Binarize \mathbf{Z} :** $\forall i \quad z_i = \arg \max_{k'} z_{ik'}$
 2. Assignment of features (8)
 • $\mathbf{W} \leftarrow \mathbf{M}_w^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5}$
 • **Binarize \mathbf{W} :** $\forall j \quad w_j = \arg \max_{k'} w_{jk'}$
until convergence of J_{RBDCo} (7)

The data co-clustering task is carried out by iteratively computing \mathbf{Z} and \mathbf{W} based on the interplay between the two updating rules derived from the maximization of the objective criterion J_{RBDCo} ,

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{W} \end{bmatrix} \leftarrow \begin{bmatrix} 0 & \mathbf{M}_z \\ \mathbf{M}_w^\top & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \\ \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \end{bmatrix} = \begin{bmatrix} \mathbf{M}_z \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \\ \mathbf{M}_w^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \end{bmatrix}. \tag{8}$$

If we set $\mathbf{M}_z = \mathbf{M}_w = \mathbf{S}_r \mathbf{X} \mathbf{S}_c$, this leads to,

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{W} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{S}_r \mathbf{X} \mathbf{S}_c \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \\ \mathbf{S}_c \mathbf{X}^\top \mathbf{S}_r \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \end{bmatrix}. \tag{9}$$

The RBDCo updating rules in (8) mutually exploit the duality of the documents and words and reinforce their joint clustering with bi-directional multiplicative regularizations using \mathbf{S}_c and \mathbf{S}_r . By generating explicit assignments of words, RBDCo produces interpretable descriptions of the resulting co-clusters. In addition, by iteratively alternating between the two updating rules, RBDCo performs an implicit adaptive word selection at each iteration and flexibly measures the distances between documents. It therefore works well with high-dimensional sparse data. The conscience mechanism embedded in RBDCo also means that it performs well with unbalanced document-term data (see Sect. 6). Algorithm 2 details the alternating procedure of RBDCo.

In the case of a *symmetric* regularization, we set $\mathbf{M}_z = \mathbf{M}_w = \mathbf{M}$, i.e., the same regularization is applied to the update rules for both \mathbf{Z} and \mathbf{W} . The objective of RBDCo is then reduced to

$$J_{RBDCo} \equiv Tr(\mathbf{Z}^\top \mathbf{M} \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5}). \tag{10}$$

If $\mathbf{M}_z = \mathbf{M}_w = \mathbf{X}$, then RBDCo is equivalent to the particular case DCC. In fact, comparing (8) and (6), it is easy to see that RBDCo generalizes DCC—DCC being RBDCo with all similarity matrices equal to \mathbf{I} —.

5.2 A generalized regularization framework

RBDCo offers a highly flexible framework in the context of text data co-clustering for the integration of supplementary information embedded in matrices that encapsulate similarities between documents and semantic correlations between words. We distinguish two types of regularization: (i) **symmetric regularization**, which consists in the application of the same regularization for the update of \mathbf{Z} and \mathbf{W} ($\mathbf{M}_z = \mathbf{M}_w$), and (ii) **asymmetric regularization**, which considers different regularizations for the update of \mathbf{Z} and \mathbf{W} ($\mathbf{M}_z \neq \mathbf{M}_w$). Table 1 summarizes the different symmetric and asymmetric configurations covered by RBDCo.

The different regularization schemes described in Table 1 highlight the flexibility of the proposed model and the connections with other approaches that can derive from it (Sect. 5.3). In our study, we indicated and justified the choice of the model retained for the case of *document-term* data on which we focused (see *Particular cases* and Sect. 6.3.1). For other types of data, the user may favor one model over another. An automatic model selection could be part of an interesting future study.

Particular cases The high degree of flexibility offered by RBDCo for multiplicative bi-directional regularizations gives rise to a variety of versions. For instance, if the identity matrix is assigned to the right-hand side of the regularization matrices \mathbf{M}_z and \mathbf{M}_w^\top , we obtain the *asymmetric uncross* case:

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{W} \end{bmatrix} \leftarrow \begin{bmatrix} 0 & \mathbf{S}_r \mathbf{X} \\ \mathbf{S}_c \mathbf{X}^\top & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \\ \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \end{bmatrix} = \begin{bmatrix} \mathbf{S}_r \mathbf{X} \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \\ \mathbf{S}_c \mathbf{X}^\top \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \end{bmatrix}. \tag{11}$$

Similarly, if the identity matrix is assigned to the left-hand side of the \mathbf{M}_z and \mathbf{M}_w^\top regularization matrices, we obtain the *asymmetric cross* case:

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{W} \end{bmatrix} \leftarrow \begin{bmatrix} 0 & \mathbf{X} \mathbf{S}_c \\ \mathbf{X}^\top \mathbf{S}_r & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \\ \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \end{bmatrix} = \begin{bmatrix} \mathbf{X} \mathbf{S}_c \tilde{\mathbf{W}} \mathbf{D}_z^{-0.5} \\ \mathbf{X}^\top \mathbf{S}_r \tilde{\mathbf{Z}} \mathbf{D}_w^{-0.5} \end{bmatrix}. \tag{12}$$

This second particular case, $RBDCo_{[S_c, S_r]}$, usually produces the best performance with document-text data (see Sect. 6). Here, row/document clustering \mathbf{Z} is regularized with the word co-occurrence information \mathbf{S}_c , and column/word clustering \mathbf{W} is regularized with the document content similarities \mathbf{S}_r . This *cross*-regularization is the most *natural* bi-directional regularization, reflecting the iterative alternating projections of words in the document space, and vice versa.

Table 1 Description of RBDCo regularization schemes

Regularization type	Data regularization		Notation
	M_z	M_w	
Symmetric	\mathbf{X}	\mathbf{X}	$\text{RBDCo}_{[I, I]}$
Symmetric	$\mathbf{S}_r \mathbf{X}$	$\mathbf{S}_r \mathbf{X}$	$\text{RBDCo}_{[S_r, S_r]}$
Symmetric	$\mathbf{X} \mathbf{S}_c$	$\mathbf{X} \mathbf{S}_c$	$\text{RBDCo}_{[S_c, S_c]}$
Symmetric	$\mathbf{S}_r \mathbf{X} \mathbf{S}_c$	$\mathbf{S}_r \mathbf{X} \mathbf{S}_c$	$\text{RBDCo}_{[S_r, S_c, S_r, S_c]}$
Asymmetric	$\mathbf{S}_r \mathbf{X}$	\mathbf{X}	$\text{RBDCo}_{[S_r, I]}$
Asymmetric	$\mathbf{X} \mathbf{S}_c$	\mathbf{X}	$\text{RBDCo}_{[S_c, I]}$
Asymmetric	\mathbf{X}	$\mathbf{S}_r \mathbf{X}$	$\text{RBDCo}_{[I, S_r]}$
Asymmetric	\mathbf{X}	$\mathbf{X} \mathbf{S}_c$	$\text{RBDCo}_{[I, S_c]}$
Asymmetric (<i>uncross</i>)	$\mathbf{S}_r \mathbf{X}$	$\mathbf{X} \mathbf{S}_c$	$\text{RBDCo}_{[S_r, S_c]}$
Asymmetric (<i>cross</i>)	$\mathbf{X} \mathbf{S}_c$	$\mathbf{S}_r \mathbf{X}$	$\text{RBDCo}_{[S_c, S_r]}$

5.3 Connection to matrix decomposition

5.3.1 Connection to NMF

Basically, $\tilde{\mathbf{Z}} = \mathbf{Z} \mathbf{D}_z^{-0.5}$ denotes the likelihood of documents being associated with document clusters, and $\tilde{\mathbf{W}} = \mathbf{W} \mathbf{D}_w^{-0.5}$ the likelihood of words being associated with word clusters. The ij th entry of $\tilde{\mathbf{Z}} \tilde{\mathbf{W}}^T$ therefore indicates the possibility that the j th word will be present in the i th-document, computed as the dot product of the i th row of $\tilde{\mathbf{Z}}$ and the j th row of $\tilde{\mathbf{W}}$. Hence, $\tilde{\mathbf{Z}} \tilde{\mathbf{W}}^T$ can be interpreted as the approximation of the original data \mathbf{X} . Our goal is then to find a \mathbf{Z} and a \mathbf{W} that minimize the squared error between \mathbf{X} and its approximation $\tilde{\mathbf{Z}} \tilde{\mathbf{W}}^T$. From a *Nonnegative Matrix Factorization* (NMF) perspective (Lee and Seung 2001), we have

$$\begin{aligned} \min_{\mathbf{Z}, \mathbf{W}} \|\mathbf{X} - \mathbf{Z} \mathbf{D}_z^{-0.5} \mathbf{D}_w^{-0.5} \mathbf{W}^T\|_F^2 &\equiv \min_{\tilde{\mathbf{Z}}, \tilde{\mathbf{W}}} \|\mathbf{X} - \tilde{\mathbf{Z}} \tilde{\mathbf{W}}^T\|_F^2 \\ &\equiv \max_{\tilde{\mathbf{Z}}, \tilde{\mathbf{W}}} \text{Tr}(\tilde{\mathbf{Z}}^T \mathbf{X} \tilde{\mathbf{W}}) \\ &\equiv \max_{\Theta, \mathbf{Z}} L_c(\Theta | \mathbf{X}, \mathbf{Z}). \end{aligned}$$

It will be remarked that, by construction, both $\tilde{\mathbf{Z}}$ and $\tilde{\mathbf{W}}^T$ are non-negative and orthogonal; we have $\tilde{\mathbf{Z}}^T \tilde{\mathbf{Z}} = \mathbf{D}_z^{-0.5} \mathbf{Z}^T \mathbf{Z} \mathbf{D}_z^{-0.5} = \mathbf{I}$, and similarly, we also have $\tilde{\mathbf{W}}^T \tilde{\mathbf{W}} = \mathbf{D}_w^{-0.5} \mathbf{W}^T \mathbf{W} \mathbf{D}_w^{-0.5} = \mathbf{I}$. Our proposed generalized regularization framework RBDCo therefore allows us to see that in its basic configuration, vMF model-based co-clustering with a conscience mechanism is in fact equivalent to a double orthogonal NMF applied to spherical data.

5.3.2 Link to NMTF

In a similar way, we can identify the link to Non-negative Matrix Tri-Factorization. Let us consider the weighting matrix $\mathbf{D} = \mathbf{D}_z^{-0.5} \mathbf{D}_w^{-0.5}$, which is diagonal by construction and where each diagonal value \mathbf{D}_{kk} represents the square root

of the geometric mean of document and word cluster sizes in block k . It follows that $\min_{\mathbf{Z}, \mathbf{W}} \|\mathbf{X} - \mathbf{Z} \mathbf{D}_z^{-0.5} \mathbf{D}_w^{-0.5} \mathbf{W}^T\|_F^2$ is equivalent to

$$\begin{aligned} \min_{\mathbf{Z}, \mathbf{W}, \mathbf{D}=\mathbf{D}_z^{-0.5} \mathbf{D}_w^{-0.5}} \|\mathbf{X} - \mathbf{Z} \mathbf{D} \mathbf{W}^T\|_F^2 &\equiv \max_{\mathbf{Z}, \mathbf{W}, \mathbf{D}} \text{Tr}(\mathbf{Z}^T \mathbf{X} \mathbf{W} \mathbf{D}) \\ &\equiv \max_{\Theta, \mathbf{Z}} L_c(\Theta | \mathbf{X}, \mathbf{Z}) \end{aligned}$$

which is also equivalent to fast NMTF proposed in (Wang et al. 2011), with an additional constraint on the centroid matrix \mathbf{D} in order to meet the requirement of directional data.

5.3.3 Link to spectral co-clustering

If, on the other hand, we relax the non-negativity constraint on both $\tilde{\mathbf{Z}}$ and $\tilde{\mathbf{W}}$, we have

$$\max_{\Theta, \mathbf{Z}} L_c(\Theta | \mathbf{X}, \mathbf{Z}) \equiv \max_{\tilde{\mathbf{Z}}^T \tilde{\mathbf{Z}}=\mathbf{I}, \tilde{\mathbf{W}}^T \tilde{\mathbf{W}}=\mathbf{I}} \text{Tr}(\tilde{\mathbf{Z}}^T \mathbf{X} \tilde{\mathbf{W}}),$$

where $\tilde{\mathbf{Z}} = \mathbf{Z} \mathbf{D}_z^{-0.5}$ and $\tilde{\mathbf{W}} = \mathbf{W} \mathbf{D}_w^{-0.5}$. It is easy to verify that $\tilde{\mathbf{Z}}$ and $\tilde{\mathbf{W}}$ satisfy the orthogonality constraint, i.e., $\tilde{\mathbf{Z}}^T \tilde{\mathbf{Z}} = \mathbf{I}$ and $\tilde{\mathbf{W}}^T \tilde{\mathbf{W}} = \mathbf{I}$. This optimization problem can be transformed using Lagrange multipliers into an eigenvalue problem. Then, given $\text{svd}(\mathbf{X}) = \tilde{\mathbf{Z}} \Sigma \tilde{\mathbf{W}}^T$, the discrete co-clustering is obtained by performing k-means on the concatenated data $[\mathbf{Z} \ \mathbf{W}]^T$. This is equivalent to the spectral co-clustering method proposed in (Dhillon and Modha 2001).

5.4 Link to block seriation

The basic idea of block co-clustering consists in modelling the simultaneous row and column partitions using a block seriation relation \mathbf{Q} defined on $I \times J$ (where I is the set of objects and J the set of attributes). Given that $\mathbf{Q} = \mathbf{Z} \mathbf{W}^T$, the general term can be expressed as follows: $q_{ij} = 1$ if object

i is in the same block as attribute j , and $\mathbf{q}_{ij} = 0$ otherwise. Thus we have

$$\mathbf{q}_{ij} = \sum_{k=1}^g \mathbf{z}_{ik} \mathbf{w}_{jk} = \left(\mathbf{Z}\mathbf{W}^\top\right)_{ij}. \tag{13}$$

The matrix \mathbf{Q} represents a block seriation relation (see (Marcotorchino 1991) for further details) that must respect the following properties:

- *Binarity* $\mathbf{q}_{ij} \in \{0, 1\}, \forall (i, j) \in I \times J$.
- *Assignment constraints* These constraints ensure the bijective correspondence between classes in two partitions, meaning that each class in the partition of I has one corresponding class in the partition of J , and vice versa. These constraints are expressed linearly as follows:

$$\begin{cases} \sum_{j \in J} \mathbf{q}_{ij} \geq 1 \quad \forall i \in I \\ \sum_{i \in I} \mathbf{q}_{ij} \geq 1 \quad \forall j \in J. \end{cases}$$

- *Triad impossible* The role of these constraints is to ensure the disjoint structure of the blocks, which is expressed by the following system inequality:

$$\begin{cases} \mathbf{q}_{ij} + \mathbf{q}_{i'j'} + \mathbf{q}_{i'j} - \mathbf{q}_{i'j'} - 1 \leq 1 \\ \mathbf{q}_{i'j'} + \mathbf{q}_{i'j} + \mathbf{q}_{ij} - \mathbf{q}_{i'j} - 1 \leq 1 \\ \mathbf{q}_{i'j} + \mathbf{q}_{ij} + \mathbf{q}_{i'j'} - \mathbf{q}_{i'j'} - 1 \leq 1 \\ \mathbf{q}_{ij} + \mathbf{q}_{i'j'} + \mathbf{q}_{i'j} - \mathbf{q}_{i'j} - 1 \leq 1. \end{cases}$$

These constraints also generalize transitivity for non-symmetric data. In the case where $I = J$, it is easy to show that the block seriation relation \mathbf{Q} becomes an equivalence relation, i.e., $\mathbf{Q} = \mathbf{Z}\mathbf{Z}^\top$ or $\mathbf{Q} = \mathbf{W}\mathbf{W}^\top$.

It will be remarked that (13) is not balanced in terms of the cluster sizes for rows and columns, meaning that a cluster might become small when affected by outliers. For this reason, we propose a new scaled block seriation relation that considers both row and column cluster sizes:

$$\tilde{\mathbf{q}}_{ij} = \sum_{k=1}^g \frac{\mathbf{z}_{ik} \mathbf{w}_{jk}}{\sqrt{\mathbf{z}_{.k} \mathbf{w}_{.k}}} = \sum_{k=1}^g \tilde{\mathbf{z}}_{ik} \tilde{\mathbf{w}}_{jk} = \left(\tilde{\mathbf{Z}}\tilde{\mathbf{W}}^\top\right)_{ij}. \tag{14}$$

A new measure, which we call *scaled block seriation criterion*, is defined as follows:

$$\begin{aligned} \min_{\mathbf{Z}, \mathbf{W}} \|\mathbf{X} - \mathbf{Z}\mathbf{D}_z^{-0.5} \mathbf{D}_w^{-0.5} \mathbf{W}^\top\|_F^2 &\equiv \min_{\tilde{\mathbf{Z}}, \tilde{\mathbf{W}}} \|\mathbf{X} - \tilde{\mathbf{Z}}\tilde{\mathbf{W}}^\top\|_F^2 \\ &\equiv \min_{\tilde{\mathbf{Q}}} \|\mathbf{X} - \tilde{\mathbf{Q}}\|_F^2 \\ &\equiv \max_{\tilde{\mathbf{Q}}} \text{Tr} \left(\mathbf{X}\tilde{\mathbf{Q}}^\top\right) \\ &\equiv \max_{\Theta, \mathbf{Z}} L_c(\Theta|\mathbf{X}, \mathbf{Z}). \end{aligned}$$

This is a scaled variant of the block seriation method (Marcotorchino 1991).

5.5 RBDCo regularization matrices

The regularization matrices \mathbf{S}_c and \mathbf{S}_r are built from the original document-term matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$. We first consider \mathbf{S}_c , for which we use a nonlinear transformation of the word co-occurrences, namely the Pointwise Mutual Information (PMI). It must be emphasized that the PMI has been shown to be strongly correlated with human assessment for word relatedness (Newman et al. 2009; Role and Nadif 2011). However, in other contexts, the user can easily introduce his own specific information about words/documents meaning. The PMI between words w_i and w_j is defined as $\log(p(w_i, w_j)/p(w_i)p(w_j))$. Assuming that the documents are the context in which words co-occur, and using the matrix $\mathbf{C} = \mathbf{X}^\top \mathbf{X}$, we can empirically estimate the PMI as follows:

$$\text{PMI}_C(w_i, w_j) = \log \frac{c_{ij} \times c_{..}}{c_{j.c.j}}, \tag{15}$$

where $c_{..} = \sum_{ij} c_{ij}$, $c_{i.} = \sum_j c_{ij}$ and $c_{.j} = \sum_i c_{ij}$. PMI values can be positive or negative. Positive values indicate that a word pair co-occurs more than by chance. Negative values are harder to interpret, since they would seem to indicate word pairs that co-occur *less* than by chance. A generally accepted approximation consists in replacing all negative values with 0, giving the Positive Pointwise Mutual Information (PPMI). One advantage of the PPMI is that it reduces the density of the PMI matrix. In RBDCo, we consider the PPMI_c matrix as our word regularization matrix \mathbf{S}_c . Similarly, we can compute a matrix \mathbf{R} such that $\mathbf{R} = \mathbf{X}\mathbf{X}^\top$. In virtue of (15), we can define a $\text{PMI}_r(d_i, d_j)$ between documents d_i and d_j that gives the co-occurrence frequency of two documents in the latent space of words. Just like in the case of \mathbf{S}_c , we consider the PPMI_R as being \mathbf{S}_r .

We have chosen to make use of PPMI for RBDCo regularization matrices, since these matrices are very general and suitable for incorporating side similarity information. They can also be computed quite easily from the original data matrix. However, other document or word embeddings obtained via external methods might also be used (e.g., Word2Vec (Mikolov et al. 2013), Doc2Vec (Le and Mikolov 2014)).

6 Experimental analysis

We now present our extensive experimental results that show the good performance of our method across a wide range of *real-world* text datasets. We first compare several variants of our RBDCo approach (Sect. 6.3.1). We then evaluate

the best RBDCo variant in relation to baseline clustering and co-clustering methods (Sect. 6.3.2). Specifically, the clustering algorithms that we consider are k -means, *spectral clustering* (Spec), Non-Negative Matrix Factorization (NMF), and *spherical k -means* (Skmeans). It is generally recognized that Skmeans in particular is well-suited to high-dimensional sparse text data. The baseline co-clustering algorithms are NMTF, DCC, and CoClustMod. The latter, CoClustMod, is a recent graph modularity-based co-clustering algorithm proposed by Ailem et al. (2016), in which CoClustMod was shown, through extensive experiments on text datasets, to outperform several other established co-clustering methods designed for the same task, including the well-known spectral co-clustering (Dhillon and Modha 2001), and ITCC (Dhillon 2003). Finally, we demonstrate the effectiveness of RBDCo in comparison with a competitive regularized text co-clustering method, WC-NMTF (Salah et al. 2018). Although WC-NMTF is a regularized co-clustering approach, it is based on an *additive unidirectional* regularization, where only word co-occurrence is used. By contrast, RBDCo proposes a *bi-directional multiplicative* regularization and embeds a *conscience* mechanism that makes it able to handle strongly unbalanced textual data. k -means, Spec and NMF come from the `scikit-learn`⁴ Python package, and Skmeans and CoClustMod from the `coclust`⁵ Python package. We implemented RBDCo, DCC and WC-NMTF in Python.

6.1 Benchmark datasets

We analyzed eight benchmark datasets widely used for document clustering purposes, namely SPORTS, TR45, LA12, CLASSIC4, CSTR, OHSCALE, PUBMED5, and CLASSIC3. Each dataset can be viewed as a contingency matrix where the coefficients x_{ij} indicate the number of occurrences of word j in document i . Together, these datasets contain a number of different challenging situations, including different degrees of cluster balance, diverse cluster sizes, and various degrees of cluster overlap.

Table 2 provides an overview of the important characteristics of the datasets sorted in increasing order of their *Balance* coefficient, which is the ratio of the smallest cluster size to the largest cluster size. As is frequently the case in document-term co-clustering, within these benchmark datasets labels are known only for the documents, and not for the words. However, given that the word partition is inherently linked to the document partition, we would expect the quality of the document clustering to be informative about the quality of the word clustering.

⁴ <https://scikit-learn.org/stable/>.

⁵ <https://pypi.python.org/pypi/coclust>.

6.2 Experimental settings and evaluation

In all our experiments, the document-term count matrix is normalized using the TF-IDF weighting scheme (*term-frequency times inverse document frequency*), as implemented in the `scikit-learn` Python package. The results are averaged over 20 different runs. For RBDCo, each run is done with 10 different initializations and a number of iterations below 100. Specifically, the final RBDCo co-clustering is automatically obtained based on the best criterion (Eq. 7) among the different initializations. To avoid poor local solutions that could be produced by early hard word assignments in the RBDCo iteration, we perform stochastic column assignments during the first 70 iterations, as described in (Salah and Nadif 2019). Whenever applicable, the approaches that RBDCo is being compared with were also performed with 10 initializations and not more than 100 iterations.

We evaluate the document clustering quality of RBDCo using two measures that are widely used for assessing the similarity between the estimated clustering and the true clustering. The measures are *Normalized Mutual Information* (NMI) (Strehl and Ghosh 2003) and *Adjusted Rand Index* (ARI) (Hubert and Arabie 1985; Steinley 2004). Specifically, NMI evaluates to what extent the estimated clustering is informative about the known clustering, and ARI quantifies the agreement between the estimated clustering and the true labels. NMI is less sensitive than ARI to cluster splitting or merging.

6.3 Empirical results on document clustering

6.3.1 Comparing RBDCo variants

We first compare the four RBDCo versions that incorporate information on both the document and the word dimensions, namely $\text{RBDCo}_{[S_c, S_r]}$, $\text{RBDCo}_{[S_r, S_c]}$, $\text{RBDCo}_{[S_c, S_c]}$ and $\text{RBDCo}_{[S_r, S_r]}$ (see Table 1 for details on RBDCo schemes). Table 3 summarizes the NMI and ARI evaluations for these versions on all the benchmark datasets.

These four RBDCo schemes give good NMI and ARI results, with a null standard deviation for almost all datasets. Overall, $\text{RBDCo}_{[S_c, S_r]}$ is the most effective variant. As already mentioned (Sect. 5.2), this version has the most *natural* bi-directional regularization for co-clustering. SPORTS is an exception: for this dataset it is the *uncross* bi-directional regularization $\text{RBDCo}_{[S_r, S_c]}$ that provides the best results (Table 3, first row). This may be explained by its high degree of document cluster imbalance. Among all the datasets, SPORTS has the lowest balance coefficient (0.036) and the lowest *ratio of minimum to expected* (RME = 0.099)—this corresponds to the smallest cluster size with respect to the expected cluster size, that is to say n/g , where g is the number of clusters. SPORTS also has the greatest *standard devia-*

Table 2 Description of Datasets

Datasets	Characteristics				
	#documents	#words	g	Sparsity (%)	Balance
SPORTS	8580	14870	7	99.14	0.036
TR45	690	8261	10	96.60	0.088
LA12	6279	31472	6	99.52	0.281
CLASSIC4	7094	5896	4	99.41	0.323
CSTR	475	1000	4	96.60	0.399
OHSCALE	11162	11465	10	99.47	0.437
PUBMED5	12648	19518	5	99.68	0.580
CLASSIC3	3891	4303	3	98.95	0.710

Table 3 Mean±sd clustering NMI and ARI on documents × terms matrices for RBDC_o variants

Datasets		RBDC _o [S _c ,S _r]	RBDC _o [S _r ,S _c]	RBDC _o [S _c ,S _c]	RBDC _o [S _r ,S _r]
SPORTS	NMI	0.67 ± 0.01	0.71 ± 0.00	0.64 ± 0.00	0.70 ± 0.00
	ARI	0.57 ± 0.03	0.68 ± 0.00	0.57 ± 0.00	0.65 ± 0.00
TR45	NMI	0.76 ± 0.00	0.74 ± 0.00	0.76 ± 0.00	0.74 ± 0.00
	ARI	0.68 ± 0.00	0.67 ± 0.00	0.68 ± 0.00	0.68 ± 0.00
LA12	NMI	0.58 ± 0.00	0.55 ± 0.00	0.56 ± 0.03	0.57 ± 0.01
	ARI	0.56 ± 0.00	0.53 ± 0.00	0.54 ± 0.03	0.53 ± 0.01
CLASSIC4	NMI	0.77 ± 0.00	0.75 ± 0.00	0.76 ± 0.00	0.76 ± 0.00
	ARI	0.78 ± 0.00	0.76 ± 0.00	0.77 ± 0.00	0.78 ± 0.00
CSTR	NMI	0.78 ± 0.00	0.73 ± 0.00	0.77 ± 0.00	0.73 ± 0.00
	ARI	0.82 ± 0.00	0.77 ± 0.00	0.81 ± 0.00	0.78 ± 0.00
OHSCALE	NMI	0.44 ± 0.00	0.44 ± 0.01	0.44 ± 0.00	0.45 ± 0.00
	ARI	0.35 ± 0.00	0.34 ± 0.01	0.34 ± 0.00	0.36 ± 0.00
PUBMED5	NMI	0.91 ± 0.00	0.90 ± 0.00	0.91 ± 0.00	0.91 ± 0.00
	ARI	0.94 ± 0.00	0.94 ± 0.00	0.94 ± 0.00	0.94 ± 0.00
CLASSIC3	NMI	0.95 ± 0.00	0.93 ± 0.00	0.95 ± 0.00	0.93 ± 0.00
	ARI	0.97 ± 0.00	0.96 ± 0.00	0.97 ± 0.00	0.96 ± 0.00

Bold values indicate the best result over all methods

tion in cluster sizes (SDCS = 1253.01), which is defined as $\{1/(g - 1) \sum_{k=1}^g (n_k - n/g)^2\}^{0.5}$, where n is the total number of documents and n_k is the cardinality of the k th cluster. This dataset therefore requires that the document similarity S_r be applied on the document dimension rather than on the word dimension in order to achieve a significantly higher NMI (0.71) and ARI (0.68).

Below, in the light of these results, we will consider only the cross-regularized RBDC_o scheme RBDC_o[S_c,S_r], where the document clustering Z is regularized with word co-occurrence information S_c , and the word clustering W is regularized with document content similarity S_r .

6.3.2 Evaluating RBDC_o against baselines

Table 4 is a synopsis of our results for RBDC_o and the other clustering/co-clustering methods, comparing their NMI and ARI values across the various benchmark datasets. The first thing to notice is that RBDC_o clearly outperforms the

standard or competitive co-clustering methods shown in the rightmost three columns of Table 4 (see also *Average ranks*). In contrast to DCC, CoClustMod and NMF, RBDC_o uses two regularization terms, specifically a word correlation matrix S_c and a document similarity matrix S_r . The superior performance of RBDC_o can therefore be attributed to these regularization terms. The ARI metric is generally more sensitive than NMI to cluster merging or splitting. However, RBDC_o has good performances for both NMI and ARI, even for highly unbalanced datasets (such as TR45 and SPORTS).

As expected, our co-clustering approach generally outperforms baseline clustering methods for the document clustering, with a higher mean margin for NMF, k -means and Spec than for the co-clustering approaches. It will be remarked that Skmeans performs well on two benchmark datasets, PubMed5 and CLASSIC3, with a small mean margin of 0.02 for NMI and 0.01 for ARI. However, our approach significantly outperforms Skmeans on unbalanced text datasets (from SPORTS to OHSCALE), with a

Table 4 RBDCo baseline comparisons with clustering and co-clustering approaches

Datasets		RBDCo _[S_c, S_r]	Skmeans	NMF	k-means	Spec	DCC	CoClustMod	NMTEF
SPORTS	NMI	0.67 ± 0.01	0.64 ± 0.02	0.53 ± 0.02	0.44 ± 0.03	0.42 ± 0.00	0.57 ± 0.01	0.55 ± 0.05	0.52 ± 0.03
	ARI	0.57 ± 0.03	0.47 ± 0.03	0.28 ± 0.03	0.14 ± 0.04	0.11 ± 0.00	0.39 ± 0.01	0.49 ± 0.06	0.28 ± 0.03
TR45	NMI	0.76 ± 0.00	0.65 ± 0.04	0.58 ± 0.05	0.63 ± 0.04	0.52 ± 0.00	0.69 ± 0.02	0.51 ± 0.04	0.63 ± 0.03
	ARI	0.68 ± 0.00	0.58 ± 0.06	0.44 ± 0.08	0.49 ± 0.09	0.32 ± 0.01	0.56 ± 0.04	0.45 ± 0.04	0.53 ± 0.04
LA12	NMI	0.58 ± 0.00	0.55 ± 0.02	0.43 ± 0.02	0.42 ± 0.04	0.35 ± 0.01	0.52 ± 0.02	0.43 ± 0.02	0.41 ± 0.02
	ARI	0.56 ± 0.00	0.50 ± 0.04	0.39 ± 0.04	0.26 ± 0.07	0.17 ± 0.02	0.45 ± 0.03	0.41 ± 0.03	0.35 ± 0.04
CLASSIC4	NMI	0.77 ± 0.00	0.69 ± 0.01	0.53 ± 0.01	0.54 ± 0.00	0.48 ± 0.00	0.72 ± 0.00	0.67 ± 0.01	0.55 ± 0.03
	ARI	0.78 ± 0.00	0.49 ± 0.00	0.44 ± 0.01	0.37 ± 0.01	0.29 ± 0.00	0.71 ± 0.01	0.61 ± 0.03	0.44 ± 0.01
CSTR	NMI	0.78 ± 0.00	0.68 ± 0.03	0.64 ± 0.01	0.62 ± 0.03	0.52 ± 0.00	0.68 ± 0.01	0.67 ± 0.04	0.67 ± 0.03
	ARI	0.82 ± 0.00	0.70 ± 0.05	0.57 ± 0.04	0.53 ± 0.06	0.35 ± 0.00	0.61 ± 0.05	0.71 ± 0.05	0.63 ± 0.09
OHSCALE	NMI	0.44 ± 0.00	0.42 ± 0.02	0.39 ± 0.02	0.37 ± 0.01	0.33 ± 0.00	0.37 ± 0.02	0.32 ± 0.02	0.40 ± 0.03
	ARI	0.35 ± 0.00	0.34 ± 0.02	0.30 ± 0.03	0.23 ± 0.03	0.22 ± 0.00	0.28 ± 0.01	0.23 ± 0.02	0.32 ± 0.03
PUBMED5	NMI	0.91 ± 0.00	0.94 ± 0.02	0.86 ± 0.05	0.74 ± 0.09	0.68 ± 0.00	0.58 ± 0.05	0.48 ± 0.07	0.81 ± 0.09
	ARI	0.94 ± 0.00	0.96 ± 0.02	0.88 ± 0.08	0.63 ± 0.16	0.53 ± 0.00	0.51 ± 0.05	0.45 ± 0.09	0.79 ± 0.13
CLASSIC3	NMI	0.95 ± 0.00	0.96 ± 0.00	0.75 ± 0.20	0.90 ± 0.01	0.62 ± 0.00	0.94 ± 0.00	0.94 ± 0.00	0.91 ± 0.00
	ARI	0.97 ± 0.00	0.98 ± 0.00	0.75 ± 0.24	0.94 ± 0.00	0.54 ± 0.00	0.97 ± 0.00	0.97 ± 0.00	0.95 ± 0.00
Mean difference	NMI	<i>reference</i>	0.04	0.14	0.15	0.24	0.10	0.16	0.12
	ARI	<i>reference</i>	0.08	0.20	0.26	0.39	0.15	0.17	0.17
Average ranks	NMI	1.25	2	5.25	5.75	7.5	3.37	5.37	4.75
	ARI	1.25	2.25	5.25	6.25	7.75	3.87	4.12	4.5

Mean ± sd clustering NMI and ARI on documents × terms matrices. Bold values indicate the best result over all methods




Fig. 2 NMI and ARI comparison of $RBDCo_{[S_c, S_c]}$, WC-NMTF and DCC on documents \times terms matrices

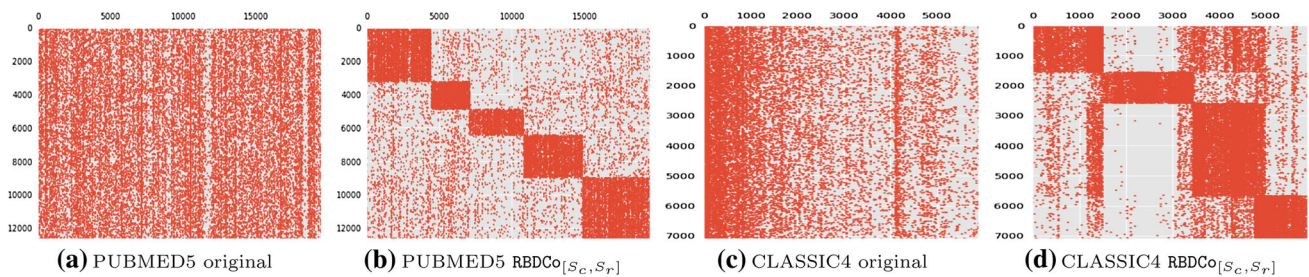


Fig. 3 Original datasets (a and c) and reorganized version (b and d) using $RBDCo_{[S_c, S_r]}$

mean margin of 0.06 for NMI and 0.11 for ARI (Table 4, first two columns). These good results on text datasets against a strong competitor like *Skmeans* are an indication of how well $RBDCo$ is able to deal with a wide range of textual datasets, and in particular in relation to text data with very small or very large (co-)clusters (see Table 4, *Average ranks*). It should also be remembered that $RBDCo$ provides a simultaneous clustering of documents *and* words and hence ensures the identification of the main document cluster topics, while *Skmeans* gives only a one-sided clustering without automatic association between documents and words.

6.3.3 Bi-directional word-based regularization

We will now compare $RBDCo$ with a competitive regularized method, namely WC-NMTF. WC-NMTF uses an additive *uni-directional* regularization, where only word co-occurrence is used on the partition of columns to obtain block diagonal co-clusters. A comparison of $RBDCo_{[S_c, S_c]}$ with WC-NMTF allows us to evaluate the advantage to be derived from multiplicative *bi-directional* word similarity regularization (i.e., S_c applied on the columns *and* the rows of X). To further enrich our comparison, we also consider the DCC evaluations. As detailed in Sect. 3, DCC, just like $RBDCo$, incorporates a conscience mechanism, but it does not use any regularization. Figure 2 gives the NMI and ARI measures for $RBDCo_{[S_c, S_c]}$, WC-NMTF, and DCC, on all datasets. These results clearly show the better performance of $RBDCo$. In particular, $RBDCo$ improves the PUBMED5 document parti-

tioning almost by a factor of two in relation to WC-NMTF, although the bi-directional regularization uses only word similarity information.

6.4 Empirical results on word clustering

6.4.1 Block diagonal co-clustering

$RBDCo$ partitions the data in diagonal document-term co-clusters, resulting in a document *and* a word clustering. Figure 3b, d shows the structures revealed by $RBDCo_{[S_c, S_r]}$ for PUBMED5 and CLASSIC4 (dots indicate strong TF-IDF weights). While the original PUBMED5 matrix does not have any explicit structure (Fig. 3a), $RBDCo$ proposes a very clear co-clustering of the documents and terms (Fig. 3b; NMI = 0.91, ARI = 0.94). Figure 3d shows the structure uncovered by $RBDCo$ for CLASSIC4 (NMI = 0.77, ARI = 0.78).

The words that occur most frequently within a co-cluster C_i are usually considered to be the most representative terms for that co-cluster. The ranking of these top terms obtained with $RBDCo$ is given in Table 5 for PUBMED5 and CLASSIC4. We can see that these terms provide a good interpretability of the partitioning. Most importantly, they can easily be linked to the topics of the true document classes. PUBMED5 is composed of five document classes, namely Age-related Macular Degeneration (AMD), Otitis, Kidney Stones, Hay Fever and Migraine. $RBDCo$ clearly uncovers associated topic words. Similarly, for CLASSIC4 $RBDCo$ gives top terms that are highly indica-

Table 5 Top frequent terms with $RBDCO_{\{s_i, s_j\}}$ and average PPMI *within* (w/i) and *between* (b/w) co-clusters

		PUBMED5					CLASSIC4				
	AMD (C1)	Otitis (C2)	Kidney Stones (C3)	Hay Fever (C4)	Migraine (C5)		CISI (C1)	CACM (C2)	MEDLINE (C3)	CRANFIELD (C4)	
	wa	otiti	stone	allerg	migrain		librari	algorithm	cell	flow	
	macular	ear	renal	nasal	patient		inform	system	patient	pressure	
	ey	children	calcium	rhiniti	headach		scienc	program	increas	number	
	visual	media	urinari	symptom	studi		research	comput	normal	boundari	
	amd	middl	kidnei	pollen	thi		studi	method	growth	layer	
	retin	acut	oxal	season	treatment		index	languag	rat	effect	
	acuiti	antibiot	rate	effect	ar		develop	gener	group	result	
	associ	aom	urin	allergen	pain		book	data	blood	heat	
	degen	tube	percutan	asthma	clinic		servic	problem	treatment	wing	
	group	om	calculi	increa	compar		retriev	time	tissu	bodi	
PPMI											
<i>w/i</i>	0.68	1.37	1.67	1.29	0.25	PPMI	0.69	0.42	1.10	0.84	
<i>b/w</i>	0.04	0.02	0.03	0.07	0.08		0.08	0.05	0.06	0.04	
<i>random</i>	0.19						0.20				

Table 6 NPMI_i scores *within* (*w/i*) and *between* (*b/w*) top frequent word clusters

	<i>w/i</i>	<i>b/w</i>	Otitis	<i>w/i</i>	<i>b/w</i>	Kidney Stones	<i>w/i</i>	<i>b/w</i>	Hay Fever	<i>w/i</i>	<i>b/w</i>	Migraine	<i>w/i</i>	<i>b/w</i>
AMD														
macular	0.67	0.07	otiti	0.48	0.06	urinari	0.55	0.24	allerg	0.55	0.22	migrain	0.54	0.16
degen	0.59	0.15	infect	0.44	0.29	excret	0.52	0.11	rhiniti	0.55	0.06	headach	0.40	0.24
retin	0.50	0.19	pneumonia	0.44	0.17	uret	0.50	-0.01	allergi	0.54	0.20	triptan	0.39	0.00
edema	0.46	0.27	bacteri	0.41	0.19	urin	0.50	0.20	asthma	0.51	0.22	treatment	0.33	0.27
diabet	0.42	0.23	acut	0.39	0.28	kidnei	0.50	0.24	allergen	0.51	0.09	pain	0.32	0.24
acuti	0.41	0.08	chronic	0.39	0.32	renal	0.49	0.26	immunotherapi	0.40	0.08	patient	0.32	0.29
visual	0.36	0.09	antibiot	0.38	0.22	uric	0.48	0.12	pollen	0.38	0.06	efficaci	0.32	0.21
amd	0.32	0.02	recurr	0.34	0.18	oxal	0.44	0.02	nasal	0.37	0.16	climic	0.31	0.23
inject	0.32	0.19	effu	0.34	0.08	acid	0.41	0.14	symptom	0.34	0.30	drug	0.29	0.22
ey	0.29	0.12	complic	0.33	0.24	metabol	0.40	0.16	eosinophil	0.32	0.03	sever	0.25	0.22

tive of the true document classes, namely CISI (information retrieval), CACM (computing machinery), MEDLINE (medical) and CRANFIELD (aeronautical systems). We recall that PUBMED5 and CLASSIC4 contain stemmed terms, so, for example, we have ‘ey’ rather than ‘eye’, and ‘studi’ rather than ‘study’.

6.4.2 Quantifying the quality of word clusters

Assessing the quality of word clustering is challenging, since the benchmark datasets commonly used in text document co-clustering provide the true document labels only. To assess the quality of the incorporation of word similarity information on the word partitioning, we first focus on the PPMI score. We then propose an enhanced version of the NPMI (Normalized Pointwise Mutual Information) score, to quantify the quality of the word clustering.

PPMI-score assessment We first consider the ten most frequent terms in each word cluster as the top terms. Table 5 gives the average pairwise PPMI values for the ten most frequent terms *within* and *between* word clusters. A *random* average is also given as a reference. This is an average pairwise PPMI over 100 random groups of ten words from the 1000 most frequent words within the whole corpus. Interestingly, for both the PUBMED5 and the CLASSIC4 datasets, the *within* PPMI average is much higher than the *between* PPMI average, indicating that RBDC₀ makes effective use of the PPMI regularization information given as input throughout its alternating iterations in a way that ultimately favors the grouping of semantically related words. The average pairwise PPMI between RBDC₀ word clusters even exhibits a stronger antagonism than the average pairwise PPMI for random groups of words ($PPMI_{random} \sim 0.2$). Word clusters with less specific vocabulary have a lower PPMI average (e.g., **Migraine** from PUBMED5: Table 5). Although informative, these average PPMI evaluations on most frequent words cannot properly assess the quality of the word clusters. One drawback of the PPMI value is that it is unbounded. Furthermore, the frequency of the words might not be the most appropriate ranking score for identifying representative words. As an example, the terms *wa* and *group* can be found among the most frequent terms in the **AMD** word cluster for PUBMED5. In addition, word clusters might also contain different subtopics that are indirectly related to each other. Therefore, considering the average PPMI over *all* pairs of words has the effect of lowering the global score, and leads to the spurious conclusion that the word cluster is not coherent.

NPMI-score assessment We propose evaluating the word cluster coherence using the Normalized PMI (NPMI) via a *k*-nn-like (*k* nearest-neighbor) approach. The NPMI ranges between -1 and +1 and is formally defined as $NPMI(w_i, w_j) = PMI(w_i, w_j) / \log(p(w_i, w_j))$. For each

word, we propose computing an $NPMI_i$ score defined as

$$NPMI_i = \frac{1}{|\Omega_i|} \sum_{w_j \in \Omega_i} NPMI(w_i, w_j) \tag{16}$$

where Ω_i is the set of k words w_j having the highest NPMI score with w_i . $NPMI_i$ quantifies the degree to which a word belongs to a cluster, based on its relationships with its k closest NPMI neighbors. $NPMI_i$ scores can be computed within or between word clusters. Table 6 gives the top $NPMI_i$ words for PUBMED5 word clusters ($k = 5$ neighbors among the 30 most frequent terms). Probabilities are derived from the whole English Wikipedia, using a NPMI implementation proposed by Röder et al. (2015). The $NPMI_i$ scores are therefore independent of the input document-term data and regularization matrices. The top $NPMI_i$ terms contain new meaningful words rather than the most frequent terms. It can be seen that *triptan*, which is a drug specific to migraine, is in third position in the **Migraine** cluster. The top **Hay Fever** terms now include *immunotherapy*, a seasonal allergy treatment, and *oesinophil*, a marker in seasonal allergic rhinitis. The word *pneumonia* can be found in the top **Otitis** terms, reflecting the fact that *Streptococcus pneumoniae* is the most common microbial agent found in otitis. Finally, the top **AMD** terms include *diabet*, an AMD risk factor, *edema*, a symptom of macular degeneration, and *inject*, which that corresponds to *intravitreal injection*, a treatment for AMD (*intravitr* is found in 14th position with $NPMI_i = 0.21$).

Figure 4 gives more insight into the relationships between the top $NPMI_i$ words in the case of **AMD**. The color of the vertices reflects the $NPMI_i$ word score, with warmer colors corresponding to higher scores, and the thickness of the edges represents the strength of the pairwise NPMI coefficient. The most important word in terms of $NPMI_i$ is *macular*. Directly related to this word are words that generally define the disease, namely *amd* and *degen* for *Age-Related Macular Degeneration*, and *neovascular* for *advanced neovascular AMD*, which is a serious type of AMD. In the upper part of the graph, we have eye-related vocabulary (e.g., *acuit(i)*, *visual*). On the right we see the words *risk*, *factor*, and *associ*, all of which are linked to *diabet*, an AMD risk factor. The *intravitr inject* bigram is directly related to *edema* and *diabet*. This makes sense, given that intravitreal injection is a treatment for diabetic macular edema. All in all, our results based on the $NPMI_i$ score and NPMI coefficient indicate that the word clusters obtained with RBDCo are highly coherent (see Sect. 6.4.3 for the remaining word clusters on PUBMED5).

Figure 5 shows the average of the top ten $NPMI_i$ terms within and between PUBMED5 word clusters obtained using RBDCo. The within average $NPMI_i$ score is seen to be higher than the between average $NPMI_i$ score. The cluster with the strongest $NPMI_i$ with other clusters is the **Migraine** word

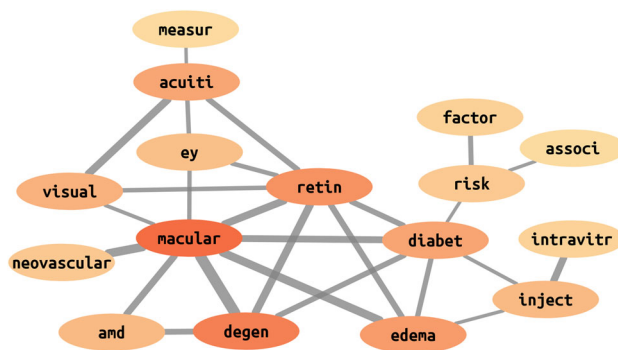


Fig. 4 NPMI graph of RBDCo AMD word cluster on PUBMED5

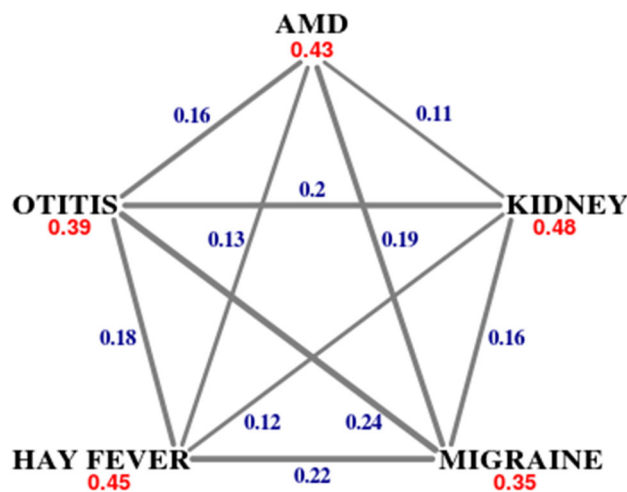


Fig. 5 Mean top $NPMI_i$ for RBDCo clusters on PUBMED5

cluster, for which the top terms are related to the topic but not specific to it.

6.4.3 RBDCo word cluster graphs on PUBMED5

We now discuss the coherence of the word cluster graphs obtained with $RBDCo_{[s_c, s_r]}$ on PUBMED5. The graphs are constructed based on the pairwise NPMI coefficient between the terms that have the highest individual $NPMI_i$ score (Eq. 16, main text). The color of the vertices reflects the $NPMI_i$ word score, with warmer colors corresponding to higher scores, and the thickness of the edges represents the strength of the pairwise NPMI value.

Otitis (Fig. 6). We observe that *otitis* is logically related to *ear* and *media*, with *otitis media* (OM) being a group of inflammatory diseases of the middle ear. The two main types are *acute otitis media* and *otitis media with effusion*. Our graph relates *acute* and *effusion* to *otitis*. Other words generally used to characterize otitis can also be found in the graph, such as *recurrent* and *chronic*. Furthermore, *S.pneumoniae* and *H.influenzae* are the most common causes of OM. *S.pneumoniae* is also the main cause of recurrent

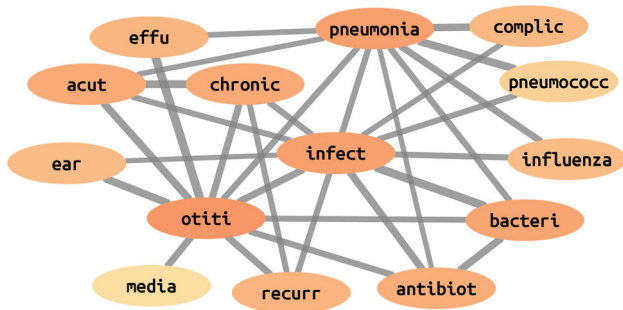


Fig. 6 NPMI graph of RBDC_O Otitis word cluster on PUBMED5




Fig. 7 NPMI graph of RBDC_O Migraine word cluster on PUBMED5

infections and postinfectious complications. This is fully coherent with the proposed graphical representation.

Migraine (Fig. 7). The most important word in this graph with respect to its NPMI_i score is *migraine*, which is strongly related to *headache* and *aura*. *Aura* is a neurological phenomenon that can accompany migraine, manifesting itself in the form of visual, sensory, and motor disturbances. As for *triptan*, this refers to a family of drugs that have been clinically assessed as being effective in the treatment of pain. We can find all these terms related in our graph, in addition to *sumatriptan*, a common migraine medication. General headache qualifiers (e.g., *severe*, *pain*) are also present.

Kidney Stones (Fig. 8). The *kidney*(y) stones or *renal calculi* graph is enriched with urinary-tract-related vocabulary, such as *ureter*, *urine*, *uric* and *urinari*(y). The graph also contains common bigrams including *urinari tract* and *uric acid*. The kidney has a clearance function that requires the excretion (*excret*) of certain metabolites (*metabol*) by our organism. In addition, the definition of kidney stones is to be seen (on the right): they are solid masses made of *crystals* that form *calcium oxalate* stones. Finally, we note the presence of *shock wave lithotripsy*(y), the most common treatment for kidney stones.

Hay fever (Fig. 9). Hay fever, also known as *allergic rhinitis*, is a *seasonal allergi*(y) caused in large part by *pollen*. These terms are related in the graph, with *pollen* linked to *allergen*. The words *intranasal* and *nasal* refer to common hay fever medications (e.g., corticosteroid nasal spray, intranasal antihistamine). The term *immunotherapy*(y), also to be seen in the

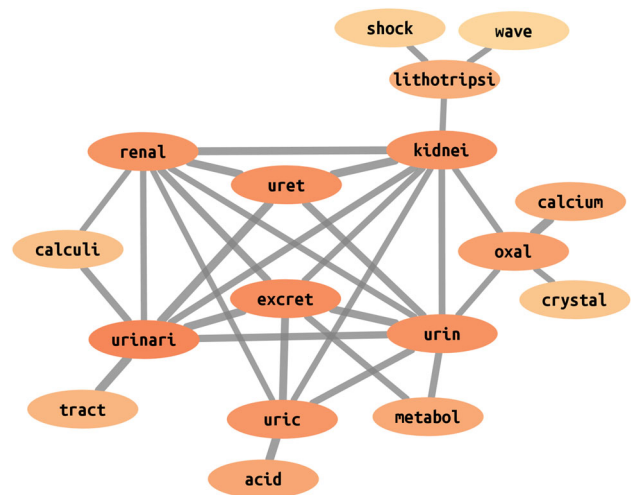


Fig. 8 NPMI graph of RBDC_O Kidney Stones word cluster on PUBMED5




Fig. 9 NPMI graph of RBDC_O Hay Fever word cluster on PUBMED5

graph, is a treatment for hay fever involving a desensitization through doses of certain allergens (e.g., grass and pollen). Our graphical representation also contains *oesinophil* and *cell*, which are linked. In fact, *oesinophil* is a specialized *cell* within the immune system that helps promote inflammation and plays a key role in the symptoms of asthma and allergies. Interestingly, the terms *allergi*, *rhiniti* and *asthma* are linked to *sympton*. In fact, the asthma and allergic rhinitis symptoms are so close that people with asthma may not recognize that they also have allergic rhinitis.

7 Conclusion

We proposed a flexible general framework, RBDC_O, for text data matrix co-clustering. RBDC_O derives from a von Mises–Fisher model-based co-clustering suitable for data that is high-dimensional, sparse, and unbalanced. Specifically, we defined our model with a matrix formulation suitable for the incorporation of complementary information to improve the co-clustering. Under some constraints, this formulation bears a close relationship to the well-known Spherical k-means, NMF and NMTF. Our approach utilizes the directional nature of text data and outperforms existing methods by

using bi-directional multiplicative regularizations to incorporate side information on the document and word dimensions. Our experiments demonstrate the good performance of RBDC_O, its robustness despite its random initialization, and its capabilities in terms of co-clustering quality and interpretability. Although all versions of RBDC_O give good results, our results suggest that RBDC_O[S_c, S_r] should be preferred. The proposed bi-directional regularization may also be seen as a means of performing a semi-supervised co-clustering. The regularization matrices might contain side expert information on the data to be partitioned, thus allowing a co-clustering that does not only depend on the input data.

In our study, we assumed that the number of co-clusters is known. Often, in practice, the number of clusters is not known and needs to be determined by the user. Assessing the number of clusters is, however, not straightforward and remains one of the biggest challenges in co-clustering. Unfortunately, in our approach we cannot rely on the well-established statistical theory of model selection, since our algorithm is not based on the maximization of the likelihood or, more precisely, on the complete-data likelihood. However, based on the vMF-Fisher mixture model designed for co-clustering, Salah and Nadif (2019) showed that (AIC) (Akaike 1998) and AIC3 (Bozdogan 2000) are effective. They also showed that these criteria give better results than the versions of the Bayesian information criterion (BIC) (Schwarz 1978) and the integrated classification likelihood (ICL) (Keribin et al. 2015) derived from *latent block models* (Govaert and Nadif 2003, 2005). As Salah and Nadif pointed out in their paper, this is due to the effective number of free parameters in the vMF-Fisher mixture model. Inspired by this result, our objective is now to address this issue in future work.

To go further, in the future we are planning to improve upon our proposed method in measuring the impact of each matrix S_r and S_c in the construction of the objective function by considering two different weights for both matrices.

Acknowledgements This work was supported by a grant overseen by the French National Research Agency (ANR) (ANR-19-CE23-0002). It also received the labelling of *Cap Digital* and *EuroBiomed/Cancer-Bio-Santé* competitiveness clusters.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

- Ahalt, S.C., Krishnamurthy, A.K., Chen, P., Melton, D.E.: Competitive learning algorithms for vector quantization. *Neural Netw.* **3**(3), 277–290 (1990)
- Ailem, M., Role, F., Nadif, M.: Graph modularity maximization as an effective method for co-clustering text data. *Knowl. Based Syst.* **109**, 160–173 (2016)
- Ailem, M., Role, F., Nadif, M.: Model-based co-clustering for the effective handling of sparse data. *Pattern Recognit.* **72**, 108–122 (2017a)
- Ailem, M., Salah, A., Nadif, M.: Non-negative matrix factorization meets word embedding. In: Proceedings of the 40th international acm sigir conference on research and development in information retrieval, pp. 1081–1084 (2017b)
- Akaike, H.: Information theory and an extension of the maximum likelihood principle. In: Selected Papers of Hirotugu Akaike, pp. 199–213. Springer (1998)
- Banerjee, A., Dhillon, I.S., Ghosh, J., Sra, S.: Clustering on the unit hypersphere using von Mises-Fisher distributions. *J. Mach. Learn. Res.* **6**, 1345–1382 (2005)
- Banerjee, A., Ghosh, J.: Frequency-sensitive competitive learning for scalable balanced clustering on high-dimensional hyperspheres. *IEEE Trans. Neural Netw.* **15**(3), 702–719 (2004)
- Bock, H.: Simultaneous clustering of objects and variables. In: Tomasone, R. (ed.) *Analyse des Données et Informatique*, pp. 187–203. INRIA, Le Chesnay (1979)
- Bock, H.-H.: Co-clustering for object by variable data matrices. In: *Advanced Studies in Behaviormetrics and Data Science*, pp. 3–17. Springer (2020)
- Bozdogan, H.: Akaike's information criterion and recent developments in information complexity. *J. Math. Psychol.* **44**(1), 62–91 (2000)
- Cho, H., Dhillon, I.S.: Coclustering of human cancer microarrays using minimum sum-squared residue Coclustering. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **5**(3), 385–400 (2008)
- Deodhar, M., Ghosh, J.: Scoal: a framework for simultaneous co-clustering and learning from complex data. *ACM Trans. Knowl. Discov. Data* **4**(3), 1–31 (2010)
- DeSieno, D.: Adding a conscience to competitive learning. In: IEEE international conference on neural networks, vol. 1, pp. 117–124, San Diego, CA, USA. Institute of Electrical and Electronics Engineers New York, IEEE (1988)
- Dhillon, I.S., Mallela, S., Modha, D.S.: Information-theoretic co-clustering. In: The 9th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 89–98 (2003)
- Dhillon, I.S., Modha, D.S.: Concept decompositions for large sparse text data using clustering. *Mach. Learn.* **42**(1–2), 143–175 (2001)
- Gopal, S., Yang, Y.: Von mises-fisher clustering models. In: International Conference on Machine Learning, pp. 154–162. PMLR (2014)
- Govaert, G.: Classification croisée. Thèse d'état, Université Paris 6, France (1983)
- Govaert, G., Nadif, M.: Clustering with block mixture models. *Pattern Recognit.* **36**(2), 463–473 (2003)
- Govaert, G., Nadif, M.: An EM algorithm for the block mixture model. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**(4), 643–647 (2005)
- Govaert, G., Nadif, M.: Block clustering with Bernoulli mixture models: comparison of different approaches. *Comput. Stat. Data Anal.* **52**(6), 3233–3245 (2008)
- Govaert, G., Nadif, M.: Co-clustering: Models, Algorithms and Applications. Wiley, New York (2013)
- Govaert, G., Nadif, M.: Mutual information, phi-squared and model-based co-clustering for contingency tables. *Adv. Data Anal. Classif.* **12**(3), 455–488 (2018)
- Hanczar, B., Nadif, M.: Ensemble methods for biclustering tasks. *Pattern Recognit.* **45**(11), 3938–3949 (2012)
- Hartigan, J.A.: Direct clustering of a data matrix. *J. Am. Stat. Assoc.* **67**(337), 123–129 (1972)
- Hofmann, T., Puzicha, J.: Latent class models for collaborative filtering. In: IJCAI, vol. 99, pp. 688–693, Stockholm, Sweden. Morgan Kaufmann (1999)

- Hubert, L., Arabie, P.: Comparing partitions. *J. Classif.* **2**(1), 193–218 (1985)
- Kerbin, C., Brault, V., Celeux, G., Govaert, G.: Estimation and selection for the latent block model on categorical data. *Stat. Comput.* **25**(6), 1201–1216 (2015)
- Le, Q., Mikolov, T.: Distributed representations of sentences and documents. *Int. Conf. Mach. Learn.* **32**, 1188–1196 (2014)
- Lee, D.D., Seung, H.S.: Algorithms for non-negative matrix factorization. *Adv. Neural Inf. Process. Syst.* **23**, 556–562 (2001)
- Madeira, S.C., Oliveira, A.L.: Biclustering algorithms for biological data analysis: a survey. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **1**(1), 24–45 (2004)
- Marcotorchino, F.: Seriation problems: an overview. *Appl. Stoch. Models Data Anal.* **7**(2), 139–151 (1991)
- Mardia, K.V., Jupp, P.E.: *Directional Statistics*, vol. 494. Wiley, New York (2009)
- McLachlan, G.J., Peel, D.: *Finite Mixture Models*. Wiley, New York (2004)
- Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: 1st international conference on learning representations, Arizona, USA. ICLR (2013)
- Newman, D., Karimi, S., Cavedon, L.: External evaluation of topic models. In: Australasian document computing symposium, IEEE (2009)
- Rocci, R., Vichi, M.: Two-mode multi-partitioning. *Comput. Stat. Data Anal.* **52**(4), 1984–2003 (2008)
- Röder, M., Both, A., Hinneburg, A.: Exploring the space of topic coherence measures. In: Proceedings of the 8th ACM international conference on Web search and data mining, Shanghai, China, pp. 399–408 (2015)
- Role, F., Morbieu, S., Nadif, M.: Coclust: a python package for co-clustering. *J. Stat. Softw. Artic.* **88**(7), 1–29 (2019)
- Role, F., Nadif, M.: Handling the impact of low frequency events on co-occurrence based measures of word similarity. In: Proceedings of the international conference on Knowledge Discovery and Information Retrieval (KDIR-2011). Scitepress, pp. 218–223 (2011)
- Salah, A., Ailem, M., Nadif, M.: Word co-occurrence regularized non-negative matrix tri-factorization for text data co-clustering. *AAAI Conf. Artif. Intell.* **32**, 3292–3299 (2018)
- Salah, A., Nadif, M.: Model-based von Mises-Fisher co-clustering with a conscience. In: Proceedings of the 2017 SIAM international conference on data mining. SIAM, pp. 246–254 (2017a)
- Salah, A., Nadif, M.: Social regularized von Mises-Fisher mixture model for item recommendation. *Data Min. Knowl. Discov.* **31**(5), 1218–1241 (2017b)
- Salah, A., Nadif, M.: Directional co-clustering. *Adv. Data Anal. Classif.* **13**(3), 591–620 (2019)
- Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**(2), 461–464 (1978)
- Steinley, D.: Properties of the Burt-adjusted rand index. *Psychol. Methods* **9**(3), 386 (2004)
- Strehl, A., Ghosh, J.: Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *J. Mach. Learn. Res.* **3**, 583–617 (2003)
- Tanay, A., Sharan, R., Shamir, R.: Biclustering algorithms: a survey. *Handb. Comput. Mol. Biol.* **9**(1–20), 122–124 (2005)
- Van Mechelen, I., Bock, H.-H., De Boeck, P.: Two-mode clustering methods: a structured overview. *Stat. Methods Med. Res.* **13**(5), 363–394 (2004)
- Vichi, M.: Double k-means clustering for simultaneous classification of objects and variables. *Adv. Classif. Data Anal.* (2001). https://doi.org/10.1007/978-3-642-59471-7_6
- Wang, H., Nie, F., Huang, H., Makedon, F.: Fast nonnegative matrix tri-factorization for large-scale data co-clustering. In: 22nd international joint conference on artificial intelligence (2011)
- Zhong, S., Ghosh, J.: Generative model-based document clustering: a comparative study. *Knowl. Inf. Syst.* **8**(3), 374–384 (2005)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.