Initializations for the Nonnegative Matrix Factorization

Amy N. Langville*
Department of Mathematics
College of Charleston
Charleston, SC 29424, USA
langvillea@cofc.edu

Carl D. Meyer†
Department of Mathematics
N. Carolina State University
Raleigh, NC 27695-8205, USA
meyer@math.ncsu.edu

Russell Albright
SAS Institute, Inc.
Enterprise Miner Research & Development
Cary, NC 27513-2414, USA
russell.albright@sas.com

ABSTRACT

The need to process and conceptualize large sparse matrices effectively and efficiently (typically via low-rank approximations) is essential for many data mining applications, including document and image analysis, recommendation systems, and gene expression analysis. The nonnegative matrix factorization (NMF) has many advantages to alternative techniques for processing such matrices, but its use comes with a caveat: the NMF must be initialized and the initialization selected is crucial to getting good solutions. It is well-known that good initializations can improve the speed and accuracy of the solutions of many NMF algorithms [43]. Add to this the fact that many NMF algorithms are sensitive with respect to the initialization of one or both NMF factors, and the impact of initializations becomes very important. In this paper, we compare the results of six initialization procedures (two standard and four new) on two alternating least squares algorithms, which we presented in [27].

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Search and Retrieval; I.7 [**Document and Text Processing**]: Miscellaneous

General Terms

nonnegative matrix factorization

Keywords

nonnegative matrix factorization, initializations, convergence, text mining, clustering

*Research supported in part by NSF CAREER-CCF-0546622.
†Research supported in part by NSF CCR-ITR-0113121 and NSF DMS 9714811.

1. INTRODUCTION

Nonnegative data are pervasive. Consider the following four important applications, each of which give rise to nonnegative data matrices.

- In document collections, documents are stored as vectors. Each element of a document vector is a count (possibly weighted) of the number of times a corresponding term appears in that document. Stacking document vectors one after the other creates a nonnegative term-by-document matrix that represents the entire document collection numerically.
- Similarly, in image collections, each image is represented by a vector, and each element of the vector corresponds to a pixel. The intensity and color of the pixel is given by a nonnegative number, thereby creating a nonnegative pixel-by-image matrix.
- For item sets or recommendation systems, the information for a purchase history of customers or ratings on a subset of items is stored in a non-negative sparse matrix.
- In gene expression analysis, gene-by-experiment matrices are formed from observing the gene sequences produced under various experimental conditions.

These are but four of the many interesting applications that create nonnegative data matrices (and tensors) [5].

Three common goals in mining information from such matrices are: (1) to automatically cluster similar items into groups, (2) to retrieve items most similar to a user's query, and (3) identify interpretable critical dimensions within the collection. For the past decade, a technique called Latent Semantic Indexing (LSI) [4], originally conceived for the information retrieval problem and later extended to more general text mining problems, was a popular means of achieving these goals. LSI uses a well-known factorization of the term-by-document matrix, thereby creating a low rank approximation of the original matrix. This factorization, the singular value decomposition (SVD) [18, 32], is a classic technique in numerical linear algebra.

The SVD is easy to compute and works well for points (1) and (2) above, but not (3). The SVD does not provide users with any interpretation of its mathematical factors or why it works so well. A common complaint from users is: *do the SVD factors reveal anything about the data collection?* Unfortunately, for the SVD, the answer to this question is no, as explained in the next section. However, an alternative

and much newer matrix factorization, known as *the nonnegative matrix factorization (NMF), allows the question to be answered affirmatively.* As a result, it can be shown that the NMF works nearly as well as the SVD on points (1) and (2), and further, can also achieve goal (3).

Most examples and applications of the NMF in this paper refer to text mining because this is the area with which we are most familiar. However, the phrase "term-by-document matrix" which we will use frequently throughout this paper can just as easily be replaced with gene-by-observation matrix, purchase-by-user matrix, etc., depending on the application area.

2. LOW RANK APPROXIMATIONS

Applications, such as text processing, data mining, and image processing, store pertinent information in a huge matrix. This matrix $\mathbf{A}$ is large, sparse, and often times nonnegative. In the last few decades, researchers realized that the data matrix could be replaced with a related matrix, of much lower rank. The low rank approximation to the data matrix $\mathbf{A}$ brought several advantages. The rank-k approximation, denoted $\mathbf{A}_k$, sometimes required less storage than $\mathbf{A}$. But most importantly, the low rank matrix seemed to give a much cleaner, more efficient representation of the relationship between data elements. The low rank approximation identified the most essential components of the data by ignoring inessential components attributed to noise, pollution, or inconsistencies. Several low rank approximations are available for a given matrix: QR, URV, SVD, SDD, PCA, ICA, NMF, CUR, etc. [24, 32, 42, 15]. In this section, we focus on two such approximations, the SVD and the NMF, that have been applied to data mining problems.

2.1 The Singular Value Decomposition

In 1991, Susan Dumais [17] used the singular value decomposition (SVD) to build a low rank approximation to the term-by-document matrix of information retrieval. In fact, to build a rank-k approximation $\mathbf{A}_k$ to the rank r term-by-document matrix $\mathbf{A}$, simply use the k most significant singular components, where $k < r$. That is,

$$\mathbf{A}_k = \sum_{i=1}^{k} \sigma_i \mathbf{u}_i \mathbf{v}_i^T = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^T,$$

where σ_i is the i^{th} singular value of $\mathbf{A}$, and $\mathbf{u}_i$ and $\mathbf{v}_i^T$ are the corresponding singular vectors [18]. The technique of replacing $\mathbf{A}$ with the truncated $\mathbf{A}_k$ is called Latent Semantic Indexing (LSI) because the low rank approximation reveals meanings and connections between documents that were hidden, or latent, in the original noisy data matrix $\mathbf{A}$.

Mathematically, the truncated SVD has one particularly appealing property: of all possible rank-k approximations, $\mathbf{A}_k$ is the best approximation in the sense that $\|\mathbf{A} - \mathbf{A}_k\|_F$ is as small as possible [4, 6]. Thus, the truncated SVD provides a nice baseline against which all other low-rank approximations can be judged for quantitative accuracy. This optimality property is also nice in practice. Algorithms for computing the k most significant singular components are fast, accurate, well-defined, and robust [2, 4, 18]. Two different algorithms will produce the same results up to roundoff error. Such uniqueness and computational robustness are comforting. Another advantage of the truncated SVD concerns building successive low rank approximations. Once $\mathbf{A}_{100}$ has been computed, no further computation is required if, for example, for sensitivity analysis or comparison purposes, other *lower* rank approximations are needed. That is, once $\mathbf{A}_{100}$ is available, then $\mathbf{A}_k$ is available for any $k \leq 100$.

LSI and the truncated SVD dominated text mining research in the 1990s [1, 3, 4, 7, 6, 8, 9, 10, 12, 17, 21, 23, 22, 30, 46, 47, 48]. However, LSI is not perfect. For instance, while it first appeared that the low rank approximation $\mathbf{A}_k$ would save storage over the original matrix $\mathbf{A}$, experiments showed that this was not the case. $\mathbf{A}$ is generally very sparse for text mining problems because only a small subset of the terms in the collection are used in any particular document. No matter how sparse the original term-by-document matrix is, the truncated SVD produces singular components that are almost always completely dense. In many cases, $\mathbf{A}_k$ can require more (sometimes much more) storage than $\mathbf{A}$.

Furthermore, $\mathbf{A}$ is always a nonnegative matrix, yet the singular components are mixed in sign. The SVD's loss of the nonnegative structure of the term-by-document matrix means that the factors of the truncated SVD provide no interpretability. To understand this statement, consider a particular document vector, say, column 1 of $\mathbf{A}$. The truncated SVD represents document 1, $\mathbf{A}_1$, as

$$\mathbf{A}_1 \approx \sigma_1 v_{11} \begin{pmatrix} \vdots \\ \mathbf{u}_1 \\ \vdots \end{pmatrix} + \sigma_2 v_{12} \begin{pmatrix} \vdots \\ \mathbf{u}_2 \\ \vdots \end{pmatrix} + \cdots + \sigma_k v_{1k} \begin{pmatrix} \vdots \\ \mathbf{u}_k \\ \vdots \end{pmatrix},$$

which reveals that document 1 is a linear combination of the singular vectors $\mathbf{u}_i$, also called the basis vectors. The scalar weight $\sigma_i v_{1i}$ represents the contribution of basis vector i in document 1. Unfortunately, the mixed signs in $\mathbf{u}_i$ and $\mathbf{v}_i$ preclude interpretation.

Clearly, the interpretability issues with the SVD's basis vectors are caused by the mixed signs in the singular vectors. Thus, researchers proposed an alternative low rank approximation that maintained the nonnegative structure of the original term-by-document matrix. As a result, the nonnegative matrix factorization (NMF) was created [29, 34]. The NMF replaces the role played by the singular value decomposition (SVD). Rather than factoring $\mathbf{A}$ as $\mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^T$, the NMF factors $\mathbf{A}$ as $\mathbf{W}_k \mathbf{H}_k$, where $\mathbf{W}_k$ and $\mathbf{H}_k$ are nonnegative.

2.2 The Nonnegative Matrix Factorization

Recently, the nonnegative matrix factorization (NMF) has been used to create a low rank approximation to $\mathbf{A}$ that contains nonnegative factors called $\mathbf{W}$ and $\mathbf{H}$. The NMF of a data matrix $\mathbf{A}$ is created by solving the following nonlinear optimization problem.

$$\begin{aligned} \min \|\mathbf{A}_{m\times n} &- \mathbf{W}_{m\times k}\mathbf{H}_{k\times n}\|_F^2, \\ s.t. \quad & \mathbf{W} \geq \mathbf{0}, \\ & \mathbf{H} \geq \mathbf{0}. \end{aligned} \tag{1}$$

The Frobenius norm is often used to measure the error between the original matrix $\mathbf{A}$ and its low rank approximation $\mathbf{WH}$, but there are other possibilities [14, 29, 33]. The rank of the approximation, k, is a parameter that must be set by the user.

The NMF is used in place of other low rank factorizations, such as the singular value decomposition (SVD) [32],

because of its two primary advantages: storage and interpretability. Due to the nonnegativity constraints, the NMF produces a so-called "additive parts-based" representation [29] of the data. One consequence of this is that the factors $\mathbf{W}$ and $\mathbf{H}$ are generally naturally sparse, thereby saving a great deal of storage when compared with the SVD's dense factors.

The NMF also has impressive benefits in terms of interpretation of its factors, which is, again, a consequence of the nonnegativity constraints. For example, consider a text processing application that requires the factorization of a term-by-document matrix $\mathbf{A}_{m\times n}$. In this case, k can be considered the number of (hidden) topics present in the document collection. In this case, $\mathbf{W}_{m\times k}$ becomes a term-by-topic matrix whose columns are the NMF basis vectors. The nonzero elements of column 1 of $\mathbf{W}$ (denoted $\mathbf{W}_1$), which is sparse and nonnegative, correspond to particular terms. By considering the highest weighted terms in this vector, one can assign a label or topic to basis vector 1. Figure 2.2 shows four basis vectors for one particular term-by-document matrix, the `medlars` dataset of medical abstracts, available at `http://www.cs.utk.edu/~lsi/`. For those familiar with the

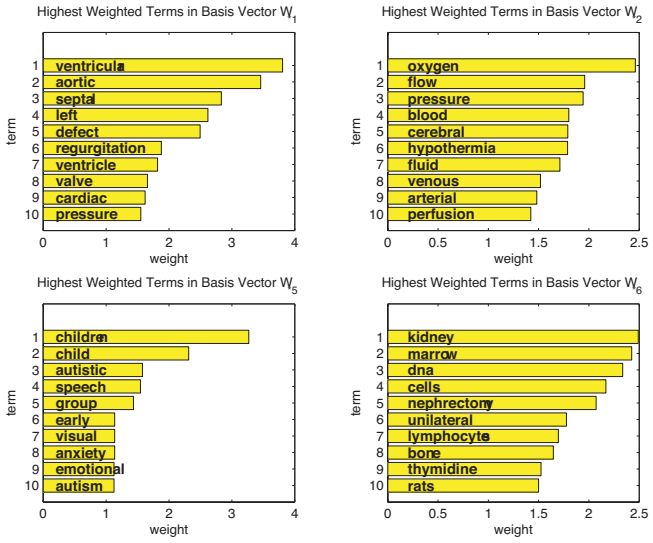

Figure 1: Interpretation of NMF basis vectors on `medlars` dataset

domain of this dataset, the NMF allows users the ability to interpret the basis vectors. For instance, a user might attach the label "heart" to basis vector $\mathbf{W}_1$ of Figure 2.2. Similar interpretation holds for the other factor $\mathbf{H}$. $\mathbf{H}_{k\times n}$ becomes a topic-by-document matrix with sparse nonnegative columns. Element j of column 1 of $\mathbf{H}$ measures the strength to which topic j appears in document 1.

Another fascinating application of the NMF is image processing. Figure 2.2 clearly demonstrates two advantages of the NMF over the SVD. First, notice that the NMF basis vectors, represented as individual blocks in the $\mathbf{W}$ matrix, are very sparse (i.e., there is much white space). Similarly, the weights, represented as individual blocks in the $\mathbf{H}_i$ vector, are also sparse. On the other hand, the SVD factors are nearly completely dense. Second, the basis vectors of the NMF, in the $\mathbf{W}$ matrix, have a nice interpretation, as individual components of the structure of the face—ears, noses, mouths, hairlines. The SVD basis vectors do not create an additive parts-based representation. In addition, the gains in storage and interpretability do not come at a loss in performance. The NMF and the SVD perform equally well in reconstructing an approximation to the original image.

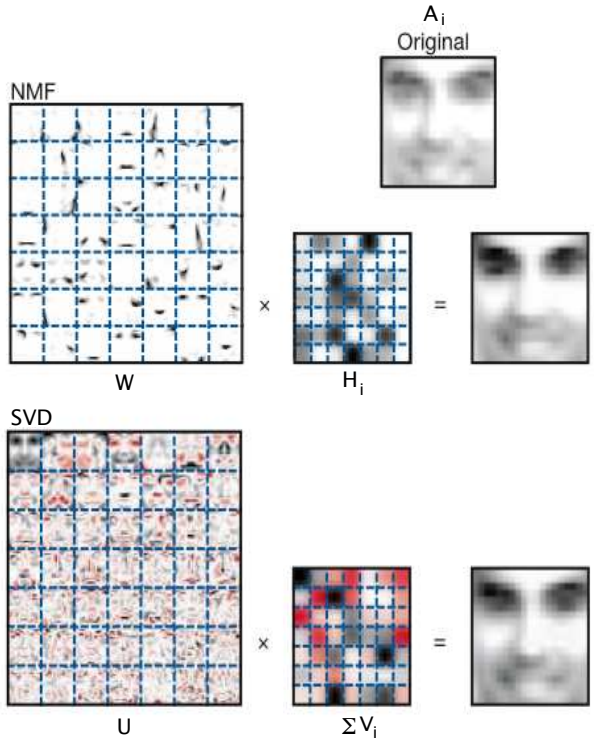

Figure 2: Interpretation of NMF and SVD basis vectors on face dataset, from [29]

Of course, the NMF has its disadvantages too. Other popular factorizations, such as the SVD, have strengths concerning uniqueness and robust computation. Yet these become problems for the NMF. There is no unique global minimum for the NMF. The optimization problem of equation (2) is convex in either $\mathbf{W}$ or $\mathbf{H}$, but not in both $\mathbf{W}$ and $\mathbf{H}$, which means that the algorithms can only, if at all, guarantee convergence to a local minimum. In practice, NMF users often compare the local minima from several different starting points, using the results of the best local minimum found. However, this is prohibitive on large, realistically-sized problems. Not only will different NMF algorithms (and there are many now [5]) produce different NMF factors, the same NMF algorithm, run with slightly different parameters, can produce different NMF factors.

2.3 Summary of SVD vs. NMF

We pause to briefly summarize the advantages of these two competing low rank approximations. The properties and advantages of the SVD include: (1) an optimality property; the truncated SVD produces the best rank-k approximation (in terms of squared distances), (2) speedy and robust compu-

tation, (3) unique factorization; initialization does not affect SVD algorithms, and (4) orthogonality; resulting basis vectors are orthogonal and allow conceptualization of original data as vectors in space. On the other hand, the advantages of NMF are: (1) sparsity and nonnegativity; the factorization maintains these properties of the original matrix, (2) reduction in storage; the factors are sparse, which also results in easier application to new data, and (3) interpretability; the basis vectors naturally correspond to conceptual properties of the data.

The strengths of one approximation become the weaknesses of another. The most severe weakness of the NMF are its convergence issues. Unlike the SVD and its unique factorization, there is no unique NMF factorization. Because different NMF algorithms can converge to different local minima (and even this convergence to local minima is not guaranteed), initialization of the algorithm becomes critical. In practice, knowledge of the application area can help guide initialization choices.

3. INITIALIZATIONS

All NMF algorithms are iterative and it is well-known that they are sensitive to the initialization of $\mathbf{W}$ and $\mathbf{H}$ [43]. Some algorithms require that both $\mathbf{W}$ and $\mathbf{H}$ be initialized [19, 20, 29, 28, 35], while others require initialization of only $\mathbf{W}$ [34, 33, 39, 40]. *In all cases, a good initialization can improve the speed and accuracy of the algorithms, as it can produce faster convergence to an improved local minimum* [42]. A good initialization can sidestep some of the convergence problems mentioned above, which is precisely why they are so important. In this section, we compare several initialization procedures (two old and four new) by testing them on the ALS algorithms presented in [27]. We choose to use the ACLS and AHCLS algorithms because they produce sparse accurate factors and require about the same time as the SVD. Most other NMF algorithms require much more time than the SVD, often times orders of magnitude more time.

3.1 Two Existing Initializations

Nearly all NMF algorithms use simple *random initialization*, i.e., $\mathbf{W}$ and $\mathbf{H}$ are initialized as dense matrices of random numbers between 0 and 1. It is well-known that random initialization does not generally provide a good first estimate for NMF algorithms [42], especially those of the ALS-type of [27] [11, 31, 36, 38]. Wild et al. [43, 44, 45] have shown that the *centroid initialization*, built from the centroid decomposition [13] is a much better alternative to random initialization. Unfortunately, this decomposition is expensive as a preprocessing step for the NMF. Because our ALS algorithms, ACLS and AHCLS, only require initialization of $\mathbf{W}$, we only discuss techniques for computing a good $\mathbf{W}^{(0)}$ (not $\mathbf{H}^{(0)}$ as well). In our algorithms, once $\mathbf{W}^{(0)}$ is known, $\mathbf{H}^{(0)}$ is computed quickly by a least squares computation.

3.2 Four New Initializations

Some text mining software produces the SVD factors for other text tasks. Thus, in the event that the SVD factor $\mathbf{V}$ is available, we propose a *SVD-centroid initialization* [26], which initializes $\mathbf{W}$ with a centroid decomposition of the low dimensional SVD factor $\mathbf{V}_{n\times k}$ [41]. While the centroid decomposition of $\mathbf{A}_{m\times n}$ can be too time-consuming, the centroid decomposition of $\mathbf{V}$ is fast because $\mathbf{V}_{n\times k}$ is much smaller than $\mathbf{A}_{m\times n}$. When the SVD factors are not available, we propose a very inexpensive procedure called *random Acol initialization*. Random Acol forms an initialization of each column of the basis matrix $\mathbf{W}$ by averaging p random columns of $\mathbf{A}$. It makes more sense to build basis vectors from the given data, the sparse document vectors themselves, than to form completely dense random basis vectors, as random initialization does. Random Acol initialization is very inexpensive, and lies between random initialization and centroid initialization in terms of performance [25, 26].

We also present two more initialization ideas, one inspired by the $\mathbf{C}$ matrix of the CUR decomposition [15], and another by the term co-occurrence matrix [37]. We refer to these last two methods as *random* $\mathbf{C}$ *initialization* and *co-occurrence initialization*, respectively. The random $\mathbf{C}$ initialization is similar to the random Acol method, except it chooses p columns at random from the longest (in the 2-norm) columns of $\mathbf{A}$, which generally means the densest columns since our text matrices are so sparse. The idea is that these might be more likely to be the centroid centers. The co-occurrence method first forms a term co-occurrence matrix $\mathbf{C} = \mathbf{A}\mathbf{A}^T$. Next, the method for forming the columns of $\mathbf{W}^{(0)}$ described as Algorithm 2 of [37] is applied to $\mathbf{C}$. The co-occurrence method is very expensive for two reasons. First, for text mining datasets, such as `reuters10`, $m >> n$, which means $\mathbf{C} = \mathbf{A}\mathbf{A}^T$ is very large and often very dense too. Second, the algorithm of [37] for finding $\mathbf{W}^{(0)}$ is extremely expensive, making this method impractical. All six initialization methods are summarized in Table 1. The two existing methods appear first, followed by our four suggested methods.

3.3 Initialization Experiments with Reuters10 dataset

The `reuters10` collection is our subset of the Reuters-21578 version of the Reuter's benchmark document collection of business newswire posts. The Reuters-21578 version contains over 20,000 documents categorized into 118 different categories, and is available online.[1] Our subset, the `reuters10` collection, is derived from the set of documents that have been classified into the top ten most frequently occurring categories. The collection contains 9248 documents from the training data of the "ModApte split" (details of the split are also available at the website above).

The numbers reported in Table 2 were generated by applying the alternating constrained least squares (ACLS) algorithm of [27] with $\lambda_H = \lambda_W = .5$ to the `reuters10` dataset. The error measure in this table is relative to the optimal rank-10 approximation given by the singular value decomposition. For this dataset, $\|\mathbf{A} - \mathbf{U}_{10}\mathbf{\Sigma}_{10}\mathbf{V}_{10}^T\|_F = 22656$. Thus, for example, the error at iteration 10 is computed as

$$\text{Error–iter.10} = \frac{\|\mathbf{A} - \mathbf{W}^{(10)}\mathbf{H}^{(10)}\|_F - 22656}{22656}.$$

We distinguish between quantitative accuracy, as reported in Table 2, and qualitative accuracy as reported in Tables 3 through 9. For text mining applications, it is often not essential that the low rank approximation be terribly precise. Often suboptimal solutions are "good enough." After

[1] `http://www.daviddlewis.com/resources/testcollections/reuters21578/`

Table 1: Initialization Methods for the NMF

Name	Proposed by	Pros	Cons
Random	Lee, Seung [28]	easy, cheap to compute	dense matrices, no intuitive basis
Centroid	Wild et al. [43]	reduces # NMF iterations, firm, intuitive foundation	expensive, must run clustering algorithm on cols of $\mathbf{A}$
SVD-Centroid	Langville [26]	inexpensive, reduces # NMF iterations	SVD factor $\mathbf{V}$ must be available
Random Acol	Langville [25]	cheap, sparse matrices built from original data	only slight decrease in number of NMF iterations
Random $\mathbf{C}$	Langville adapts from Drineas [15]	cheap, sparse	not very effective
Co-occurrence	Langville adapts from Sandler [37]	uses term-term similarities	large, dense co-occurrence matrix, very expensive computation

Table 2: Experiments with Initialization Methods for the NMF

Method	Time $\mathbf{W}^{(0)}$	Storage $\mathbf{W}^{(0)}$	Error–iter.0	Error–iter.10	Error–iter.20	Error–iter.30
Random	.09 sec	726K	4.28%	.278%	.146%	.146%
Centroid	27.72	46K	2.02%	.269%	.181%	.177%
SVD-Centroid	.65†	56K	2.08%	.057%	.057%	.057%
Random Acol*	.05	6K	2.01%	.212%	.155%	.146%
Random $\mathbf{C}^{\circ}$	.11	22K	3.35%	.287%	.199%	.189%
Co-occurrence	3287	45K	3.38%	.371%	.269%	.252%
ACLS time			.37 sec	3.42	6.78	10.29

$\dagger$ provided $\mathbf{V}$ of the SVD is already available
* each column of $\mathbf{W}^{(0)}$ formed by averaging 20 random columns of $\mathbf{A}$
$^{\circ}$ each column of $\mathbf{W}^{(0)}$ formed by averaging 20 of the longest columns of $\mathbf{A}$

reviewing Tables 3–9, it is easy to see why some initializations give better accuracy and converge more quickly. They start with basis vectors in $\mathbf{W}^{(0)}$ that are much closer to the best basis vectors found, as reported in Table 3, which was generated by using the basis vectors associated with the best global minimum for the `reuters10` dataset, found by using 500 random restarts. In fact, the relative error for this global minimum is .009%, showing remarkable closeness to the optimal rank-10 approximation. By comparing each subsequent table with Table 3, it's clear why one initialization method is better than another. The best method, SVD-centroid initialization, starts with basis vectors very close to the "optimal" basis vectors of Table 3. On the other hand, random and random Acol initialization are truly random. Nevertheless, random Acol does maintain one clear advantage over random initialization as it creates a very sparse $\mathbf{W}^{(0)}$. The Random C and co-occurrence initializations suffer from lack of diversity. Many of the longest documents in the `reuters10` collection appear to be on similar topics, thus, not allowing $\mathbf{W}^{(0)}$ to cover many of the reuters topics.

Because the algorithms did not produce the "wheat" vector always in column one of $\mathbf{W}$, we have reordered the resulting basis vectors in order to make comparisons easier. We also note that the nonnegative matrix factorization did produce basis vectors that cover 8 of the 10 "correct" reuters classifications, which appear on the last line of Table 3. The two missing reuters classifications are `corn` and `grain`, both of which are lumped into the first basis vector labeled `wheat`. This first basis vector does break into two separate vectors, one pertaining to `wheat and grain` and another to `corn` when the number of basis vectors is increased from $k = 10$ to $k = 12$. We note that these categories have been notoriously difficult to classify, as previously reported in [16].

4. FUTURE WORK

Surveying the tables in this paper, prompts us to propose another initialization as future work. The SVD-centroid initialization is the best of the six methods studied, yet it requires that the $\mathbf{V}_k$ matrix of the truncated SVD be available. If $\mathbf{V}_k$ is not available, then one would spend as much time getting this matrix by computing a truncated SVD, as he would computing the NMF. Clearly, this is an unreasonable amount of preprocessing time. Thus, in the future, we will experiment with a method for approximating $\mathbf{V}_k$ by computing the truncated SVD of a random sample of columns of $\mathbf{A}$. We suspect this would be a preprocessing step whose slight expense is worth the effort.

A common issue for many factorization algorithms is updating. Once the data collection has changed and the $\mathbf{A}$ matrix updated, what procedures exist, beyond total recomputation, for updating the factorization? This is an unstudied issue for the NMF. It is tempting to use NMF factors from the original matrix to as initializations for the updated matrix. We plan to experiment with these ideas in a subsequent paper.

5. CONCLUSION

This paper presents four new initialization techniques for the nonnegative matrix factorization. Only two, the SVD-centroid and the random Acol initialization techniques, prove beneficial. Comparing these two new initializations with the

Table 3: Basis vectors of $\mathbf{W}^{(50)}$ from *Best Global Minimum* found for `reuters10`

$\mathbf{W}_1^{(50)}$	$\mathbf{W}_2^{(50)}$	$\mathbf{W}_3^{(50)}$	$\mathbf{W}_4^{(50)}$	$\mathbf{W}_5^{(50)}$	$\mathbf{W}_6^{(50)}$	$\mathbf{W}_7^{(50)}$	$\mathbf{W}_8^{(50)}$	$\mathbf{W}_9^{(50)}$	$\mathbf{W}_{10}^{(50)}$
tonne	billion	share	stg	mln-mln	gulf	dollar	oil	loss	trade
wheat	year	offer	bank	cts	iran	rate	opec	profit	japan
grain	earn	company	money	mln	attack	curr.	barrel	oper	japanese
crop	qrtr	stock	bill	shr	iranian	bank	bpd	exclude	tariff
corn	rise	sharehol.	market	net	ship	yen	crude	net	import
agricul.	pct	common	england	avg	tanker	monetary	price	dlrs	reagan
`wheat`	`earn`	`acquisition`		`interest`	`ship`	`frgn-exch.`	`oil`		`trade`

Table 4: Basis vectors of $\mathbf{W}^{(0)}$ created by *Random Initialization* for `reuters10`

$\mathbf{W}_1^{(0)}$	$\mathbf{W}_2^{(0)}$	$\mathbf{W}_3^{(0)}$	$\mathbf{W}_4^{(0)}$	$\mathbf{W}_5^{(0)}$	$\mathbf{W}_6^{(0)}$	$\mathbf{W}_7^{(0)}$	$\mathbf{W}_8^{(0)}$	$\mathbf{W}_9^{(0)}$	$\mathbf{W}_{10}^{(0)}$
announce	wpp	formality	bulletin	matthews	dramatic	squibb	wag	cochran	erik
medtec	reflagging	simply	awfully	nyt	boca raton	kuwaiti	oils	mln	support
pac	kwik	moonie	blair	barrel	clever	dacca	hears	barriers	sale oil
purina	tilbury	tmg	fresno	purina	billion	democrat	bwtr	deluxe	direct
mezzanine	capacitor	bushnell	farm	june	bkne	induce	nestle	mkc	wheat
foreign	grain	country	leutwiler	trend	clever	rate	federal	economic	aid

Table 5: Basis vectors of $\mathbf{W}^{(0)}$ created by *Centroid Initialization* for `reuters10`

$\mathbf{W}_1^{(0)}$	$\mathbf{W}_2^{(0)}$	$\mathbf{W}_3^{(0)}$	$\mathbf{W}_4^{(0)}$	$\mathbf{W}_5^{(0)}$	$\mathbf{W}_6^{(0)}$	$\mathbf{W}_7^{(0)}$	$\mathbf{W}_8^{(0)}$	$\mathbf{W}_9^{(0)}$	$\mathbf{W}_{10}^{(0)}$
tonne	bank	share	medar	cts	iran	rate	oil	stg	strike
wheat	rate	company	mdxr	mmln	gulf	dollar	trade	bill	port
grain	dollar	offer	mlx	loss	attack	bank	price	take-up	union
corn	billion	pct	mlxx	net	iranian	currency	barrel	drain	seaman
crop	pct	stock	mich	shr	missile	market	japan	mature	worker
agriculture	trade	dlrs	troy	dlrs	tanker	monetary	opec	money	ship

Table 6: Basis vectors of $\mathbf{W}^{(0)}$ created by *SVD-Centroid Initialization* for `reuters10`

$\mathbf{W}_1^{(0)}$	$\mathbf{W}_2^{(0)}$	$\mathbf{W}_3^{(0)}$	$\mathbf{W}_4^{(0)}$	$\mathbf{W}_5^{(0)}$	$\mathbf{W}_6^{(0)}$	$\mathbf{W}_7^{(0)}$	$\mathbf{W}_8^{(0)}$	$\mathbf{W}_9^{(0)}$	$\mathbf{W}_{10}^{(0)}$
tonne	billion	share	bank	cts	iran	dollar	oil	loss	trade
wheat	year	offer	money	shr	gulf	rate	barrel	oper	japan
grain	earn	company	rate	mln	attack	curr.	opec	profit	japanese
corn	qrtr	stock	stg	net	iranian	yen	crude	cts	tariff
crop	rise	pct	market	mln-mln	missile	japan	bpd	mln	import
agricul.	pct	common	pct	rev	ship	economic	price	net	country

Table 7: Basis vectors of $\mathbf{W}^{(0)}$ created by *Random Acol Initialization* for `reuters10`

$\mathbf{W}_1^{(0)}$	$\mathbf{W}_2^{(0)}$	$\mathbf{W}_3^{(0)}$	$\mathbf{W}_4^{(0)}$	$\mathbf{W}_5^{(0)}$	$\mathbf{W}_6^{(0)}$	$\mathbf{W}_7^{(0)}$	$\mathbf{W}_8^{(0)}$	$\mathbf{W}_9^{(0)}$	$\mathbf{W}_{10}^{(0)}$
mln	fee	agl	mln	mark	loss	official	dlrs	bank	trade
denman	mortgage	tmoc	dlrs	mannesmann	mln	piedmont	oper	bancaire	viermetz
dlrs	billion	bank	share	dividend	cts	dollar	billion	austral	mln
ecuador	winley	pct	seipp	mln	maki	interest	loss	neworld	nwa
venezuela	mln	company	billion	dieter	name	tokyo	texaco	datron	cts
revenue	fed	maki	dome	gpu	kato	japanese	pennzoil	share	builder

Table 8: Basis vectors of $\mathbf{W}^{(0)}$ created by *Random* C *Initialization* for `reuters10`

$\mathbf{W}_1^{(0)}$	$\mathbf{W}_2^{(0)}$	$\mathbf{W}_3^{(0)}$	$\mathbf{W}_4^{(0)}$	$\mathbf{W}_5^{(0)}$	$\mathbf{W}_6^{(0)}$	$\mathbf{W}_7^{(0)}$	$\mathbf{W}_8^{(0)}$	$\mathbf{W}_9^{(0)}$	$\mathbf{W}_{10}^{(0)}$
analyst	dollar	economic	bank	market	analyst	analyst	analyst	trade	rate
lawson	rate	policy	rate	bank	market	industry	bank	dollar	trade
market	economist	pct	market	analyst	trade	price	currency	japan	official
trade	mark	cost	currency	price	pct	market	japan	price	bank
sterling	bank	growth	dollar	mark	last	believe	billion	japanese	market
dollar	rise	trade	trade	good	official	last	cut	pct	economist

Table 9: Basis vectors of $\mathbf{W}^{(0)}$ created by *Co-occurrence Initialization* for `reuters10`

$\mathbf{W}_1^{(0)}$	$\mathbf{W}_2^{(0)}$	$\mathbf{W}_3^{(0)}$	$\mathbf{W}_4^{(0)}$	$\mathbf{W}_5^{(0)}$	$\mathbf{W}_6^{(0)}$	$\mathbf{W}_7^{(0)}$	$\mathbf{W}_8^{(0)}$	$\mathbf{W}_9^{(0)}$	$\mathbf{W}_{10}^{(0)}$
dept.	average	agricul.	national	farmer	rate-x	aver price	plywood	wash.	trade
wheat	pct	wheat	bank	rate-x	natl	average	aqtn	trade	japan
agricul.	rate	tonne	rate	natl	avge	price	aequitron	japan	billion
tonne	price	grain	pct	avge	farmer	yield	medical	official	market
usda	billion	farm	oil	cwt	cwt	billion	enzon	reagan	japanese
corn	oil	dept.	gov.	wheat	wheat	bill	enzon	pct	import

two most popular existing initializations, we find that the two new ones require less storage than the standard default random initialization and converge to more accurate local minimums than the existing centroid initialization.

6. ADDITIONAL AUTHORS

Additional authors: James Cox (SAS Institute, Inc., email: `james.cox@sas.com`) and David Duling (SAS Institute, Inc., email: `david.duling@sas.com`).

7. REFERENCES

[1] Ricardo Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. ACM Press, New York, 1999.

[2] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, Philadelphia, 2nd edition, 1994.

[3] Michael W. Berry, editor. *Computational Information Retrieval*. SIAM, Philadelphia, 2001.

[4] Michael W. Berry and Murray Browne. *Understanding Search Engines: Mathematical Modeling and Text Retrieval*. SIAM, Philadelphia, 2nd edition, 2005.

[5] Michael W. Berry, Murray Browne, Amy N. Langville, V. Paul Pauca, and Robert J. Plemmons. Algorithms and applications for the nonnegative matrix factorization. 2006. submitted.

[6] Michael W. Berry, Zlatko Drmac, and Elizabeth R. Jessup. Matrices, vector spaces and information retrieval. *SIAM Review*, 41:335–62, 1999.

[7] Michael W. Berry and R. D. Fierro. Low-rank orthogonal decompositions for information retrieval applications. *Journal of Numerical Linear Algebra with Applications*, 1(1):1–27, 1996.

[8] Michael W. Berry and Gavin W. O'Brien. Using linear algebra for intelligent information retrieval. *SIAM Review*, 37:573–595, 1998.

[9] Katarina Blom. *Information retrieval using the singular value decomposition and Krylov subspaces*. PhD thesis, University of Chalmers, January 1999.

[10] Katarina Blom and Axel Ruhe. Information retrieval using very short Krylov sequences. In *Computational Information Retrieval*, pages 41–56, 2001.

[11] Donald S. Burdick, Xin M. Tu, Linda B. McGown, and David W. Millican. Resolution of multicomponent fluorescent mixtures by analysis of the excitation-emission-frequency array. *Journal of Chemometrics*, 4:15–28, 1990.

[12] Anirban Dasgupta, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Variable latent semantic indexing. In *Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM Press, 2005.

[13] Inderjit S. Dhillon. Concept decompositions for large sparse text data using clustering. *Machine Learning*, 42(1/2):143–175, 2001.

[14] Inderjit S. Dhillon and Suvrit Sra. Generalized nonnegative matrix approximations with Bregman divergences. In *Proceeding of the Neural Information Processing Systems (NIPS) Conference*, Vancouver, B.C., 2005.

[15] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. Fast Monte Carlo algorithms for matrices III: Computing a compressed approximate matrix decomposition. *SIAM Journal on Computing*, 2006. to appear.

[16] Susan Dumais, John Platt, David Heckerman, and Mehran Sahami. Inductive learning algorithms and representations for text categorization. In *CIKM '98: Proceedings of the seventh international conference on Information and knowledge management*, pages 148–55. ACM Press, 1998.

[17] Susan T. Dumais. Improving the retrieval of information from external sources. *Behavior Research Methods, Instruments and Computers*, 23:229–236, 1991.

[18] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 1996.

[19] Patrik O. Hoyer. Non-negative sparse coding. In *Neural Networks for Signal Processing XII (Proc. IEEE Workshop on Neural Networks for Signal Processing)*, pages 557–565, 2002.

[20] Patrik O. Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5:1457–1469, 2004.

[21] M. K. Hughey and Michael W. Berry. Improved query matching using kd-trees, a latent semantic indexing enhancement. *Information Retrieval*, 2:287–302, 2000.

[22] Eric P. Jiang and Michael W. Berry. Solving total least squares problems in information retrieval. *Linear Algebra and its Applications*, 316:137–156, 2000.

[23] Fan Jiang and Michael L. Littman. Approximate dimension equalization in vector-based information retrieval. In *The Seventeenth International Conference on Machine Learning*, pages 423–430, 2000.

[24] Tamara G. Kolda and Dianne P. O'Leary. A semi-discrete matrix decomposition for latent semantic indexing in information retrieval. *ACM Transactions on Information Systems*, 16:322–346, 1998.

[25] Amy N. Langville. Algorithms for the nonnegative matrix factorization in text mining, April 2005. Slides from SAS Meeting.

[26] Amy N. Langville. Experiments with the nonnegative matrix factorization and the reuters10 dataset, February 2005. Slides from SAS Meeting.

[27] Amy N. Langville, Carl D. Meyer, Russell Albright, James Cox, and David Duling. Alternating least squares algorithms for the nonnegative matrix factorization. 2006. preprint.

[28] D. Lee and H. Seung. Algorithms for Non-Negative Matrix Factorization. *Advances in Neural Information Processing Systems*, 13:556–562, 2001.

[29] Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.

[30] Todd A. Letsche and Michael W. Berry. Large-scale information retrieval with LSI. *Informatics and Computer Science*, pages 105–137, 1997.

[31] Shousong Li and Paul J. Gemperline. Eliminating complex eigenvectors and eigenvalues in multiway analyses using the direct trilinear decomposition method. *Journal of Chemometrics*, 7:77–88, 1993.

[32] Carl D. Meyer. *Matrix Analysis and Applied Linear Algebra*. SIAM, Philadelphia, 2000.

[33] Pentti Paatero. Least squares formulation of robust non-negative factor analysis. *Chemometrics and Intelligent Laboratory Systems*, 37:23–35, 1997.

[34] Pentti Paatero and U. Tapper. Positive matrix factorization: a non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5:111–126, 1994.

[35] V. Paul Pauca, Jon Piper, and Robert J. Plemmons. Nonnegative matrix factorization for spectral data analysis. 2005.

[36] E. Sanchez and B. R. Kowalski. Tensorial resolution: A direct trilinear decomposition. *Journal of Chemometrics*, 4:29–45, 1990.

[37] Mark Sandler. On the use of linear programming for unsupervised text classification. In *The Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Chicago, IL, 2005.

[38] R. Sands and Forrest W. Young. Component models for three-way data: an alternating least squares algorithm with optimal scaling features. *Psychometrika*, 45:39–67, 1980.

[39] Farial Shahnaz. A clustering method based on nonnegative matrix factorization for text mining. Master's thesis, University of Tennessee, Knoxville, 2004.

[40] Farial Shahnaz, Michael W. Berry, V.Paul Pauca, and Robert J. Plemmons. Document Clustering Using Nonnegative Matrix Factorization. *Information Processing & Management*, 42(2):373–386, 2006.

[41] David B. Skillicorn, S. M. McConnell, and E.Y. Soong. Handbooks of data ming using matrix decompositions. 2003.

[42] Age Smilde, Rasmus Bro, and Paul Geladi. *Multi-way Analysis*. Wiley, West Sussex, England, 2004.

[43] Stefan Wild. Seeding non-negative matrix factorizations with spherical k-means clustering. Master's thesis, University of Colorado, 2003.

[44] Stefan Wild, James Curry, and Anne Dougherty. Motivating non-negative matrix factorizations. In *Eighth SIAM Conference on Applied Linear Algebra*, Philadelphia, 2003. SIAM.

[45] Stefan Wild, James Curry, and Anne Dougherty. Improving non-negative matrix factorizations through structured initialization. *Journal of Pattern Recognition*, 37(11):2217–2232, 2004.

[46] Dian I. Witter and Michael W. Berry. Downdating the latent semantic indexing model for conceptual information retrieval. *The Computer Journal*, 41(1):589–601, 1998.

[47] Hongyuan Zha, Osni Marques, and Horst D. Simon. A subspace-based model for information retrieval with applications in latent semantic indexing. *Lecture Notes in Computer Science*, 1457:29–42, 1998.

[48] Xiaoyan Zhang, Michael W. Berry, and Padma Raghavan. Level search schemes for information filtering and retrieval. *Information Processing and Management*, 37:313–34, 2001.