

# Iterative Residual Rescaling: An Analysis and Generalization of LSI

Rie Kubota Ando and Lillian Lee  
 Department of Computer Science  
 Cornell University  
 Ithaca, NY 14853-7501  
 {kubotar, llee}@cs.cornell.edu

## ABSTRACT

We consider the problem of creating document representations in which inter-document similarity measurements correspond to semantic similarity. We first present a novel *subspace-based* framework for formalizing this task. Using this framework, we derive a new analysis of *Latent Semantic Indexing* (LSI), showing a precise relationship between its performance and the *uniformity* of the underlying distribution of documents over topics. This analysis helps explain the improvements gained by Ando's (2000) *Iterative Residual Rescaling* (IRR) algorithm: IRR can compensate for distributional non-uniformity. A further benefit of our framework is that it provides a well-motivated, effective method for automatically determining the rescaling factor IRR depends on, leading to further improvements. A series of experiments over various settings and with several evaluation metrics validates our claims.

## 1. INTRODUCTION

**Background** The rapid increase in the availability of electronic documents has created high demand for automated text analysis technologies such as document clustering, summarization, and indexing. Representations enabling accurate measurement of semantic similarities between documents would greatly facilitate such technologies. In this paper, we focus on representations in which vector directionality is used to represent a document's semantics, by which we mean its (human-interpretable) constituent concepts. This goal is to be accomplished without access to concept labels, since they are typically not available in many applications.

The vector space model (VSM) is a classic method for constructing such vector-based representations. It encodes a document collection by a *term-document matrix* whose  $[i, j]$ th element indicates the association between the  $i$ th term and  $j$ th document. However, as has been pointed out previously, VSM does not always represent semantic relatedness well; for instance, documents that do not share terms

are mapped to orthogonal vectors even if they are clearly related.

*Latent Semantic Indexing (LSI)* [9, 11] attempts to overcome this and other shortcomings by computing an approximation to the original term-document matrix; this is equivalent to *projecting* the term-document matrix onto a lower-dimensional subspace. LSI has been successfully applied to information retrieval and many language analysis tasks (e.g. [12, 13, 5, 14, 21]), prompting several studies to explain its effectiveness (e.g. [4, 20, 10, 16, 3]).

Ando [1] introduced an alternative subspace-projection method, which we call *Iterative Residual Rescaling* (IRR), that outperforms LSI by counteracting its tendency to ignore minority-class documents. This is done by repeatedly rescaling vectors to amplify the presence of documents poorly represented in previous iterations. However, Ando presented only heuristic arguments to explain IRR's success.

**Contributions** In this paper, we use the notion of subspace projection to formalize the document representation problem. Based on this framework, we provide a new theoretical analysis that shows a precise relationship between the performance of LSI and the *uniformity* of the underlying distribution of documents over topics. As a consequence, we provide an explanation for IRR's success: the rescaling it performs compensates for non-uniformities in the topic-document distribution. Moreover, our framework yields a new way to automatically adjust the amount of rescaling by estimating the non-uniformity.

To support our theoretical results, we present performance measurements both on document sets in which the topic-document distributions were carefully controlled, and on unrestricted datasets as would be found in application settings. In all cases and for all metrics, the results confirm our theoretical predictions. For instance, IRR combined with our new parameter selection technique achieved up to 10.1% higher *kappa average precision* than LSI and up to 8.7% better document clustering performance. The experiments as a whole provide strong evidence for the usefulness of our framework in general and the effectiveness of our augmented IRR in particular.

**Notational conventions** A bold uppercase letter (e.g.  $\mathbf{M}$ ) denotes a matrix; the corresponding bold lowercase letter with subscript  $i$  (e.g.  $\mathbf{m}_i$ ) denotes the matrix's  $i$ th column vector. We use  $\text{range}(\mathbf{M})$  to denote  $\mathbf{M}$ 's range, or column space:  $\{\mathbf{y} \mid \exists \mathbf{x} \text{ such that } \mathbf{y} = \mathbf{M}\mathbf{x}\}$ . When a document collection has been specified,  $n$  denotes the number of documents the collection contains.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGIR '01, September 9-12, New Orleans, Louisiana, USA.  
 Copyright 2001 ACM 1-58113-331-6/01/0009 ...\$5.00.

## 2. ANALYZING LSI

### 2.1 Topic-based similarities

Our framework for analyzing Latent Semantic Indexing and other subspace-based methods revolves around the notion of topic-based similarity. Fix an  $n$ -document collection  $\mathcal{C}$  and corresponding  $m$ -by- $n$  term-document matrix  $\mathbf{D}$ . We assume that there exists a set, denoted  $\text{topics}(\mathcal{C})$ , of  $k < n$  topics underlying  $\mathcal{C}$ . We also assume that for each topic  $t$  and document  $d$  there exists a (real-valued) *relevance score*  $\text{rel}(t, d)$ , suitably normalized so that for each  $d$ , we have  $\sum_{t \in \text{topics}(\mathcal{C})} \text{rel}(t, d)^2 = 1$ . We then define the *true topic-based similarity* between two documents  $d$  and  $d'$  as:

$$\text{sim}(d, d') = \sum_{t \in \text{topics}(\mathcal{C})} \text{rel}(t, d) \text{rel}(t, d').$$

It is convenient to summarize these similarities in a single  $n$ -by- $n$  matrix  $\mathbf{S}$ , where  $\mathbf{S}[d, d'] = \text{sim}(d, d')$ .

Note that although we assume the existence of underlying topics as the basis for the true document similarities, in contrast to other analyses (e.g. [16, 3, 10]) we do *not* assume that there is an underlying generative or probabilistic model that *creates* the term-document matrix  $\mathbf{D}$ .

### 2.2 The optimum subspace

We formulate the ultimate goal of subspace-based algorithms, such as LSI, as choosing some subspace such that projecting  $\mathbf{D}$  onto this subspace creates new document vectors whose measured similarities (i.e., cosines) more closely correspond to the true topic-based similarities.

More formally, for any subspace  $\mathcal{X} \subseteq \mathbb{R}^m$ , the unique (*orthogonal*) projection of a vector  $\mathbf{x} \in \mathbb{R}^m$  onto  $\mathcal{X}$  is given by  $\mathbf{P}_{\mathcal{X}}(\mathbf{x}) = \mathbf{B}\mathbf{B}^T\mathbf{x}$  for any  $\mathbf{B}$  whose columns form an orthonormal basis for  $\mathcal{X}$ . We define the projection of  $\mathbf{D}$  onto  $\mathcal{X}$  as the matrix  $[\mathbf{P}_{\mathcal{X}}(\mathbf{d}_1) \cdots \mathbf{P}_{\mathcal{X}}(\mathbf{d}_n)]$ , i.e., the result of projecting each of the term-document vectors. Hence, after projection onto  $\mathcal{X}$ , the similarity between the  $i$ th and  $j$ th documents is measured by

$$\cos(\mathbf{P}_{\mathcal{X}}(\mathbf{d}_i), \mathbf{P}_{\mathcal{X}}(\mathbf{d}_j)) = \frac{\mathbf{P}_{\mathcal{X}}(\mathbf{d}_i)^T \mathbf{P}_{\mathcal{X}}(\mathbf{d}_j)}{\|\mathbf{P}_{\mathcal{X}}(\mathbf{d}_i)\| \|\mathbf{P}_{\mathcal{X}}(\mathbf{d}_j)\|}.$$

The *document representation problem* is as follows: given  $\mathbf{D}$  — but *not*  $\mathbf{S}$  or even any knowledge of what the underlying topics are — find a subspace  $\mathcal{X}$  such that the entries of the *deviation matrix*<sup>1</sup>

$$\text{diff}_{\mathbf{S}, \mathbf{D}}(\mathcal{X}) = \mathbf{S} - \mathbf{P}_{\mathcal{X}}(\mathbf{D})^T \mathbf{P}_{\mathcal{X}}(\mathbf{D})$$

are small. The *optimum subspace*

$$\mathcal{X}_{\text{opt}} = \arg \min_{\mathcal{X} \subseteq \text{range}(\mathbf{D})} \|\text{diff}_{\mathbf{S}, \mathbf{D}}(\mathcal{X})\|_2,$$

with ties broken by smallest dimensionality and then arbitrarily, serves as the standard for comparison in our analysis. We denote the corresponding projection operator by  $\mathbf{P}_{\text{opt}}$ , and use  $\epsilon_{\text{opt}}$  to denote the *optimum error*  $\|\text{diff}_{\mathbf{S}, \mathbf{D}}(\mathcal{X}_{\text{opt}})\|_2$ .

<sup>1</sup>It suffices to consider the inner products  $\mathbf{P}_{\mathcal{X}}(\mathbf{d}_i)^T \mathbf{P}_{\mathcal{X}}(\mathbf{d}_j)$  rather than the cosines because, as shown in [2], if there exists  $\epsilon < 1$  that upper-bounds the magnitudes of the deviation matrix's entries, then for any  $\mathbf{d}_i$  and  $\mathbf{d}_j$ ,

$$\frac{\text{sim}(i, j) - \epsilon}{1 + \epsilon} \leq \cos(\mathbf{P}_{\mathcal{X}}(\mathbf{d}_i), \mathbf{P}_{\mathcal{X}}(\mathbf{d}_j)) \leq \frac{\text{sim}(i, j) + \epsilon}{1 - \epsilon}.$$

Note that  $\epsilon_{\text{opt}}$  need not be zero, as it may be impossible to project the given term-document matrix in such a way as to perfectly recover the true document similarities.

### 2.3 The singular value decomposition and LSI

In this section, we first briefly introduce the *singular value decomposition* (SVD) [15], since singular values are necessary for our analysis. Then, we describe LSI, which is based on the SVD.

The SVD factors an arbitrary rank- $h$  matrix  $\mathbf{Z} \in \mathbb{R}^{r \times s}$  into the following product:

$$\mathbf{Z} = \mathbf{U}\Sigma\mathbf{V}^T,$$

where the columns of  $\mathbf{U} \in \mathbb{R}^{r \times h}$  and  $\mathbf{V} \in \mathbb{R}^{s \times h}$  are orthonormal,  $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_h)$  is diagonal (following convention, we assume  $\sigma_1 \geq \dots \geq \sigma_h$ ), and the  $\sigma_i$ 's are all positive. The quantities  $\sigma_i$ ,  $\mathbf{u}_i$  and  $\mathbf{v}_i$  are called the  $i$ th *singular value*, *left singular vector*, and *right singular vector*, respectively. The left singular vectors span  $\mathbf{Z}$ 's range, and  $\sigma_1 = \|\mathbf{Z}\|_2$ .

Zeroing out all but the  $\ell < h$  largest singular values yields the least-squares optimal rank- $\ell$  approximation to  $\mathbf{Z}$ . *Latent Semantic Indexing* (LSI) [9, 11] applies this rank- $\ell$  approximation to the term-document matrix  $\mathbf{D}$ , which corresponds to projecting  $\mathbf{D}$  onto the *rank- $\ell$  LSI subspace* spanned by  $\mathbf{u}_1, \dots, \mathbf{u}_{\ell}$ :

$$[\mathbf{u}_1 \dots \mathbf{u}_{\ell}] [\mathbf{u}_1 \dots \mathbf{u}_{\ell}]^T \mathbf{D} = \mathbf{U} \underbrace{\text{diag}(\sigma_1, \dots, \sigma_{\ell}, 0, \dots, 0)}_{h \text{ entries}} \mathbf{V}^T.$$

Note that the fact that this matrix approximates  $\mathbf{D}$  well does not imply that it represents the true document similarities well, as we shall see.

Further intuition may be gained on the left singular vectors by the following observation. Let  $\text{proj}^{(j)}(\mathbf{d}_i)$  be the projection of  $\mathbf{d}_i$  onto the span of  $\mathbf{u}_1, \dots, \mathbf{u}_j$ , and let  $\mathbf{r}_i^{(j)}$  be the *residual vector*  $\mathbf{d}_i - \text{proj}^{(j-1)}(\mathbf{d}_i)$ . Then,  $\mathbf{u}_j$  is the unit vector that maximizes the following quantity:

$$g^{(j)}(\mathbf{x}) = \sum_{i=1}^n \left( |\mathbf{r}_i^{(j)}| \cos(\mathbf{r}_i^{(j)}, \mathbf{x}) \right)^2.$$

In a sense,  $\mathbf{u}_j$  resembles a weighted average of residual vectors, where longer residuals receive greater weight. Hence, the  $\ell$  left singular vectors may be thought of as representing the  $\ell$  major directions in the document collection.

### 2.4 Non-uniformity and LSI

We now state our results relating the non-uniformity of the underlying topic-document distribution to the quality of the document representation spaces derived using LSI. The main outline of our argument is to first show relations between certain singular values and certain quantities linked to our topic model, and then show how the distance between the LSI-subspace and the optimal subspace relates to these singular values. Proofs of these results, which make use of invariant subspace perturbation theorems [8, 19, 15], are sketched in the appendix of this paper and given in full in [2].

Recall that we are dealing with a fixed document collection  $\mathcal{C}$  with  $k$  underlying topics. Throughout, we use  $h$  to denote the dimensionality of the optimum subspace. For clarity, we will abuse notation by writing “ $x \in y \pm z$ ” as shorthand for “ $x \in [y - z, y + z]$ ”.

A crucial quantity in our analysis is the *dominance*  $\Delta_t$  of a given topic  $t$ :

$$\Delta_t = \sqrt{\sum_{d \in \mathcal{C}} \text{rel}(t, d)^2}.$$

(It may be helpful to observe that in the special “single-topic documents” case, where each document is relevant to only one topic in  $\text{topics}(\mathcal{C})$ , squaring  $\Delta_t$  gives exactly the number of documents in  $\mathcal{C}$  that are relevant to topic  $t$ .) We assume without loss of generality that  $\Delta_1 \geq \Delta_2 \geq \dots \geq \Delta_k$ , and for convenience set  $\Delta_i = 0$  if  $i > k$ .

Now we show that the projection of  $\mathbf{D}$  onto the optimum subspace  $\mathcal{X}_{opt}$  in some sense reveals the topic dominances. It is intuitively clear, however, that the extent to which this holds should depend to some degree both on the optimum error  $\epsilon_{opt}$  and on the *topic mingling*  $\mu(\mathcal{C}) = \left(\sum_{t, t' \in \text{topics}(\mathcal{C}), t \neq t'} \left(\sum_{d \in \mathcal{C}} \text{rel}(t, d)\text{rel}(t', d)\right)^2\right)^{1/2}$  (note that in the single-topic documents case,  $\mu(\mathcal{C}) = 0$ ). Certainly, if the optimum error is high, then we cannot expect the optimum subspace to fully reveal the topic dominances; also, if there is high topic mingling in the collection, then the topics will be fairly difficult to distinguish.

**THEOREM 2.1.** *Let  $\tau_i$  be the  $i$ th largest singular value of  $\mathbf{P}_{opt}(\mathbf{D})$ . Then,  $\tau_i^2 \in \Delta_i^2 \pm (\epsilon_{opt} + \mu(\mathcal{C}))$ .*

This result gives us leave to define  $\hat{\Delta}_{max} = \tau_1$  and  $\hat{\Delta}_{min} = \tau_h$ , where  $h$  is the dimension of  $\mathcal{X}_{opt}$ . The ratio  $\hat{\Delta}_{max}/\hat{\Delta}_{min}$  then serves as a measure of the *non-uniformity* of the topic-document distribution underlying the collection: the more the largest topic dominates the collection, the higher this ratio will tend to be.

Now we are in a position to present our main theorem. This result bounds the distance between the optimum subspace  $\mathcal{X}_{opt}$  and the same-dimensionality LSI subspace  $\mathcal{X}_{LSI}$  by a function of the non-uniformity of  $\mathcal{C}$ ’s topic-document distribution. The bound also incorporates a certain value (defined precisely in the appendix)  $\hat{\epsilon}_{vsm} \in \epsilon_{vsm} \pm \epsilon_{opt}$ , where  $\epsilon_{vsm} = \|\text{diffs}_{\mathbf{D}}(\mathcal{X}_{VSM})\|_2 = \|\mathbf{S} - \mathbf{D}^T \mathbf{D}\|_2$  is the *input error*: intuitively, if a “bad” term-document matrix is received as input, one cannot expect LSI to do well.

**THEOREM 2.2.** *Let  $\mathcal{X}_{LSI}$  be the  $h$ -dimensional LSI subspace spanned by the the first  $h$  left singular vectors of  $\mathbf{D}$ . If  $\hat{\Delta}_{min} > \sqrt{\hat{\epsilon}_{vsm}}$ , then*

$$\|\tan(\Theta(\mathcal{X}_{LSI}, \mathcal{X}_{opt}))\|_2 \leq \frac{\hat{\Delta}_{max}}{\hat{\Delta}_{min}} \cdot \frac{\sqrt{\hat{\epsilon}_{vsm}}/\hat{\Delta}_{min}}{1 - (\sqrt{\hat{\epsilon}_{vsm}}/\hat{\Delta}_{min})^2},$$

where  $\Theta$  is the canonical angle matrix measuring the distance between subspaces [8, 19].

Intuitively, what this means is that  $\mathcal{X}_{LSI}$  must be close to  $\mathcal{X}_{opt}$  when the topic-document distribution is relatively uniform and the input error is small in comparison to the  $h$ th largest topic dominance. Conversely, if the input error is fixed, our bound on LSI’s performance weakens when the underlying topic-document distribution is highly non-uniform. Finally, we note that the condition on  $\hat{\epsilon}_{vsm}$  is natural: roughly speaking, if the input error is large enough to “swamp” the dominance of the  $h$ th largest topic, then intuitively we cannot expect good results.

Finally, we note that a related result can be proved which, roughly speaking, links lower bounds on the distance between the two subspaces to non-uniformity and the input error; however, this theorem is quite technical in nature and thus is omitted. We refer the reader to [2] for details.

## 2.5 Related work: theoretical analyses of LSI

As noted above, there have been several studies analyzing LSI, using approaches such as Bayesian regression models [20] and Gaussian models [10]. Here, we concentrate on describing the work most similar in spirit to ours.

Zha et al. [22] propose a subspace-based model for LSI. Their work focuses on dimensionality selection and implementation issues regarding more accurate updating schemes. Bartell et al. [4] show that LSI can be regarded as a solution to the special Multidimensional Scaling (MDS) problem of preserving the inner products of the original document vectors. However, as noted above, this is not the same as recovering hidden topic-based similarities, especially in the case of noisy data.

Perhaps the work most similar to ours is that of Papadimitriou et al. [16] and Azar et al. [3], both of which propose to explain LSI’s success with analyses that, like ours, employ invariant subspace perturbation theorems. Papadimitriou et al. start with a probabilistic corpus model. By assuming low input error and certain conditions on singular values (which, from our perspective, can be considered to be roughly equivalent to assuming relative uniformity, although [16] did not explicitly make this connection), they show that LSI will work well with high probability. But their results are based on a *pure* probabilistic corpus model in which all the documents are single-topic and topics have associated *primary* (distinguishing) disjoint sets of terms. Thus, their analysis holds only for a very restricted class of document collections. Similarly, Azar et al. also start with particular conditions and a specialized underlying generative model to show that LSI works well for “good” documents with high probability. In contrast to these approaches, our analysis does not assume a model of term-document matrix creation, and so applies to *arbitrary* term-document matrices, with the non-uniformity and the input matrix’s quality being explicit terms in our bound.

## 3. IRR: OVERCOMING NON-UNIFORMITY

Our results from Section 2.4 indicate that we could improve the performance of LSI if we could somehow “smooth” the topic-document distribution (that is, effectively lower  $\hat{\Delta}_{max}/\hat{\Delta}_{min}$ ). We show that the *Iterative Residual Rescaling* (IRR) algorithm, introduced (but not named) and heuristically motivated by Ando [1], accomplishes this task *without prior knowledge of the assignments of documents to topics*.

### 3.1 Ando’s IRR algorithm

Recall from Section 2.3 that the left singular vectors  $\mathbf{u}_1, \mathbf{u}_2, \dots$  produced by LSI can be derived, one by one, via the following computation:

$$\mathbf{u}_j = \arg \max_{\|\mathbf{x}\|=1} \sum_{i=1}^n \left( |\mathbf{r}_i^{(j)}| \cos(\mathbf{r}_i^{(j)}, \mathbf{x}) \right)^2,$$

where the  $\mathbf{r}_i^{(j)}$  are the residuals  $\mathbf{d}_i - \text{proj}^{(j-1)}(\mathbf{d}_i)$ . Unfortunately, inspection of this formula shows that when the




Figure 1: Effect of nonuniformity on LSI, and how IRR compensates.

```

IRR( $q, \ell$ ):
   $\mathbf{R} := \mathbf{D}$  /* initialize residuals by given  $m$ -by- $n$  term-document matrix */
  For  $j := 1, 2, \dots, \ell$  /* create  $\ell$  basis vectors */
    For  $i := 1, 2, \dots, n$ 
       $\hat{\mathbf{r}}_i := |\mathbf{r}_i|^q \mathbf{r}_i$  /* rescale residuals */
       $\mathbf{b}_j := \arg \max_{\mathbf{x}: \|\mathbf{x}\|=1} \left( \sum_{i=1}^n (|\hat{\mathbf{r}}_i| \cos(\hat{\mathbf{r}}_i, \mathbf{x}))^2 \right)$ 
    For  $i := 1, 2, \dots, n$ 
      Subtract from  $\mathbf{r}_i$  its projection onto  $\mathbf{b}_j$  /* recompute residuals */
   $\mathbf{D}_{\text{IRR}} := \mathbf{B}\mathbf{B}^T\mathbf{D}$ , where  $\mathbf{B} = [\mathbf{b}_1 \ \mathbf{b}_2 \ \dots \ \mathbf{b}_\ell]$  /* new document representation */

```

Figure 2: High-level pseudocode for Ando’s IRR (i.e. before augmentation with AUTO-SCALE).

topic-document distribution is highly non-uniform, the cumulative influence of a large number of (small) residuals for a major topic can cause smaller topics to be ignored, as depicted in Figures 1(a) and (b).

Our explanation of IRR’s effectiveness is that by amplifying the length differences among residual vectors, IRR boosts the influence of minority-topic documents, thus *compensating for non-uniform topic distributions* (Figure 1(c)). More precisely, IRR, like (our formulation of) the SVD, computes basis vectors by successive maximizations, as shown in the pseudocode in Figure 2. Crucially, though, IRR’s objective function,  $\hat{g}^{(j)}$ , incorporates a scaling factor  $q$  via the scaling function  $\text{pow}(\mathbf{r}, q) = |\mathbf{r}|^q \mathbf{r}$ :

$$\hat{g}^{(j)}(\mathbf{x}) = \sum_{i=1}^n \left( |\text{pow}(\mathbf{r}_i^{(j)}, q)| \cos(\text{pow}(\mathbf{r}_i^{(j)}, q), \mathbf{x}) \right)^2.$$

This is maximized by the first left singular vector of

$$\hat{\mathbf{R}}^{(j)} = \left[ \text{pow}(\mathbf{r}_1^{(j)}, q) \ \dots \ \text{pow}(\mathbf{r}_n^{(j)}, q) \right].$$

That is, IRR *rescales* each residual vector  $\mathbf{r}_i$  at each basis vector computation, increasing the contrast between long and short residuals when  $q > 0$ . LSI is the special case in which  $q = 0$ .

### 3.2 The AUTO-SCALE method

Our discussion above argues that the degree of rescaling should depend on the uniformity of the topic-document distribution. Ando [1] did not explicitly make this connection, and hence could not take advantage of it:  $q$  was determined simply by training on held-out data. In contrast, our novel analysis allows us to exploit this connection to develop an effective estimation method — *automatic scaling factor determination* (AUTO-SCALE) — that approximates the topic-

document non-uniformity without prior knowledge of the underlying topics.

AUTO-SCALE is based on the observation that we can use the quantity  $\sum_{t \in \text{topics}(\mathcal{C})} \Delta_t^4 / n^2$  as a measure of the non-uniformity of the topic-document distribution. Of course, we don’t have access to the topic dominances  $\Delta_t$ , but we can approximate this measure by <sup>2</sup>

$$f(\mathbf{D}) \stackrel{\text{def}}{=} \left( \frac{\|\mathbf{D}^T \mathbf{D}\|_F}{n} \right)^2.$$

This approximation follows from first assuming that the input matrix is fairly good, so that  $\|\mathbf{D}^T \mathbf{D}\|_F^2$  is roughly equal to  $\|\mathbf{S}\|_F^2 = \sum_{d, d' \in \mathcal{C}} \left( \sum_{t \in \text{topics}(\mathcal{C})} \text{rel}(t, d) \text{rel}(t, d') \right)^2$ . The latter can be rewritten, after some algebra, as the quantity  $\sum_{t \in \text{topics}(\mathcal{C})} \Delta_t^4 + \mu(\mathcal{C})^4$ , which is roughly  $\sum_{t \in \text{topics}(\mathcal{C})} \Delta_t^4$  if we assume approximately single-topic documents. In practice, we set  $q$  to a linear function of  $f(\mathbf{D})$ ; this is discussed in Section 5.1.

Although the above assumptions are rather coarse, AUTO-SCALE yields good empirical results: see Sections 5 and 6.

### 3.3 Dimensionality selection

IRR’s second parameter is  $\ell$ , the dimensionality of the created subspace. One way to set this parameter is to train it on held-out data. Following Ando [1], we found that learning thresholds on the *residual ratio*  $\|\mathbf{R}^{(j)}\|_F^2 / n$  as a stopping criterion is effective for both LSI and IRR. Intuitively, this ratio describes how much is left out of the proposed subspace (of course, we do not want to reproduce the term-document matrix exactly — hence the threshold). Note that this training method allows some flexibility in the chosen di-

<sup>2</sup>The *Frobenius norm*  $\|\mathbf{X}\|_F$  is defined as  $\sqrt{\sum_{i,j} \mathbf{X}[i, j]^2}$ .

mensionality: for different data sets, the same residual ratio threshold may result in selecting a different  $\ell$ .

While training on held-out data is reasonable and is commonly employed in practice, it is a relatively expensive process. A speedier alternative arises in settings in which  $k$ , the number of topics, is pre-specified — examples include cases where the topic set is a fixed class such as the TREC topic labels, or where the application allows the user to specify the appropriate level of granularity for his or her needs. In such settings, we could simply set the dimensionality equal to  $k$  as a matter of convenience. (Indeed, [16] show that under certain strong assumptions on the data, rank- $k$  LSI should perform well. See also [2].) We describe experiments with both selection methods below.

## 4. EVALUATION METRICS

**Kappa average precision** Our first evaluation metric is the *pair-wise average precision* [1], adapted from the average precision measure commonly used in information retrieval. The motivation behind this metric is that the measured similarity for any two *intra-topic* documents (i.e., that share at least one topic) should be higher than for any two *cross-topic* documents which have no topics in common. More formally, let  $p_i$  denote the document pair with the  $i$ th largest measured similarity (cosine). Precision for an *intra-topic* pair  $p_j$  is defined by

$$\text{prec}(p_j) = \frac{\# \text{ of intra-topic pairs } p_i \text{ such that } i \leq j}{j}.$$

The *pair-wise average precision* is the average of these precision values over all intra-topic pairs.

To compensate for the effect of large topics (which increase the likelihood of chance intra-topic pairs), we modify the pair-wise average precision to create a new metric, which we call the *kappa precision* in reference to the Kappa statistic [17, 7]:

$$\text{prec}_\kappa(p_i) = \frac{\text{prec}(p_i) - \text{chance}}{1 - \text{chance}},$$

where  $\text{chance} = (\# \text{ of intra-topic pairs}) / (\# \text{ of document pairs})$ . The *kappa average precision*  $\kappa$  is defined to be the average of the kappa precision over all intra-topic pairs, and is a linear function of the pair-wise average precision.

**Clustering** We also test how well the new subspaces represent document similarities by seeing whether document clustering improves when these new representations are used as input. To simplify the scoring, we consider only single-topic documents.

Let  $C$  be a cluster-topic contingency table such that  $C[i, j]$  is the number of documents in cluster  $i$  that are relevant to topic  $j$ , as in [18]. We define  $s(C) = \sum_{i,j} \mathcal{N}_{ij} / n$ , where  $\mathcal{N}_{ij} = C[i, j]$  if  $C[i, j]$  is the *unique* maximum in both its row and column, and  $\mathcal{N}_{ij} = 0$  otherwise. Note that this (rather strict) measure only considers the most tightly coupled topic-cluster assignment, and decreases when either cluster purity or topic integrity declines (see Figure 3).

To factor out the idiosyncracies of particular clustering algorithms, we apply six standard clustering methods — single-link, complete-link, group average, and k-means with initial clusters generated by these three methods — to the document vectors in each proposed subspace, and record both the *ceiling* (highest) and *floor* (lowest)  $s(C)$  scores. While the ceiling performance is perhaps more intuitive, we

	topic 1	topic 2	topic 3	topic 4
cluster 1	5	10	<b>20</b>	0
cluster 2	5	10	5	0
cluster 3	0	0	0	<b>21</b>
cluster 4	<b>15</b>	5	0	0
cluster 5	0	0	0	4

**Figure 3: Sample contingency table, with  $s(C) = (15 + 20 + 21)/100 = 56\%$ .**

observe that the floor performance also gives us important information about the quality of the representation being evaluated: if the floor is low, then there is at least one clustering algorithm for which the document subspace is not a good representation; otherwise, the representation is good for *all six* clustering algorithms.

## 5. CONTROLLED DISTRIBUTIONS

Our first suite of experiments studies the dependence of LSI and IRR on increasingly less uniform topic-document distributions. The results strongly support our theoretical analysis of LSI’s sensitivity to non-uniformity.

### 5.1 Experimental setting

To focus on distributional non-uniformity, we first chose two TREC topics, and then specified seven *distribution types*: (25, 25), (30, 20), (35, 15), (40, 10), (43, 7), (45, 5), and (46, 4), where  $(n_1, n_2)$  indicates that  $n_i$  of the documents are relevant to topic  $i$ . For each of these types, we generated ten sets of 50 TREC documents each, where each document was relevant to exactly one of the pre-selected topics. We also created five-topic<sup>3</sup> data sets in the same manner, using distribution types of the form  $(i, j, j, j, j)$  (which makes uniformity comparisons obvious).

To create the term-document matrices, we extracted single-word stemmed terms using TALENT [6], removed stop-words, and then length-normalized the document vectors (so that term weights were frequency-based).

To implement AUTO-SCALE, we set  $q = \alpha \cdot f(\mathbf{D}) + \beta$ , where  $\alpha = 3.5$  and  $\beta = 0$  for *all* our experiments. These values (which are necessary to determine the “units” of the scale factor) were empirically determined once and for all from observations on data disjoint from our test sets. This contrasts with training  $q$  for every new test set encountered, as in [1]. Training is an expensive process, and we envision interactive applications such as organizing query results (a task we simulate in Section 6) in which what would serve as training data is not obvious. We thus view AUTO-SCALE as a practical alternative to the usual parameter training.

For simplicity, the dimensionality of LSI and IRR in our experiments was set to the number of topics.<sup>4</sup>

### 5.2 Controlled-distribution results

We first examine the kappa average precision results, shown in Figure 4. The  $x$ -axis represents the nonuniformity of the topic-document distribution, as measured by  $\Delta_{max}/\Delta_{min}$ .

<sup>3</sup>Results for three- and four-topic document sets were similar and are therefore omitted.

<sup>4</sup>Our preliminary experiments with dimensionality training indicated that indeed this was often the best dimensionality for LSI and almost always the best for AUTO-SCALE-IRR.




Figure 4: Kappa average performance, two topics. Points are averages over ten document sets.

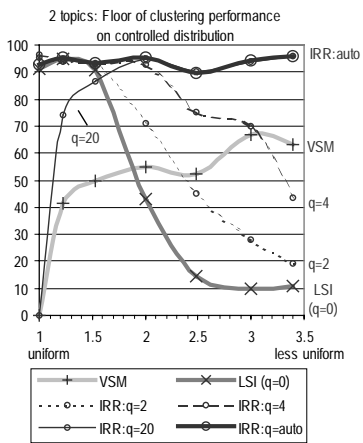


Figure 6: Floor and ceiling clustering results, two topics. Points are averages over ten document sets.





Figure 5: Automatically-determined scaling factor values: ten-set average and standard deviation.

We see that when the topic-document distribution is relatively uniform, LSI’s performance is higher than 90%. However, as the nonuniformity increases, the performance of LSI drops precipitously, in accordance with our theorems above.

Also, our interpretation of the scaling factor  $q$  as compensating for non-uniformity is borne out nicely. For highly uniform distributions, the performance difference between  $q = 0$  (at which  $IRR = LSI$ ),  $q = 2$ , and  $q = 4$  is not great. At medium nonuniformity,  $q = 0$  degrades, but  $q = 2$  still does about the same as  $q = 4$ . But as the non-uniformity increases even more, we see that  $q = 2$  is not large enough to compensate, and so declines in comparison to  $q = 4$ .

Furthermore, we see that IRR with AUTO-SCALE (labelled ‘IRR:  $q = \text{auto}$ ’) does extremely well across all levels of non-uniformity. Figure 5 shows that AUTO-SCALE indeed adjusts for more non-uniform distributions: the chosen scaling factor increases on average as the non-uniformity goes up.

Now, one might conjecture that instead of using AUTO-SCALE, it would suffice simply to choose a single very large value of  $q$ . Intuitively, though, this is problematic, since too high a scaling factor would tend to completely eliminate residuals. Furthermore, the  $q = 20$  curve in Figure 4 disproves the conjecture: in the uniform case, selecting

an overly large scaling factor hurts performance, driving it below the baseline VSM curve.

The two-topic floor and ceiling clustering results, shown in Figure 6, exhibit precisely the same types of behaviors as in the kappa average precision case. The floor performances are especially interesting, as they show that AUTO-SCALE-IRR exhibits very good performance for all six of our rather wide variety of clustering algorithms. They also indicate that VSM is ‘fragile’ for uniform distributions, in that sometimes it is a very poor representation for at least one of the clustering algorithms we employed.

Finally, Figure 7 shows the results of the same evaluation experiments run on five-topic data. Again, the empirical results are completely in line with what we predicted, with AUTO-SCALE leading to strong performance over all metrics and all degrees of non-uniformity. Note that the gap between LSI and VSM decreases in comparison to the  $k = 2$  case; this is due to the fact that at higher dimensionalities, the subspace produced by LSI gets closer to that of the original term-document matrix.

These results all strongly support our theoretical claims.

## 6. UNRESTRICTED DISTRIBUTIONS

In this section, we experiment on the more realistic setting of document sets without distribution restrictions. We expect that in practice, topic-document distributions will be fairly non-uniform, so that IRR should perform well in comparison to LSI. Figure 8 summarizes the evaluation settings.

We used 648 TREC documents, each relevant to exactly one of twenty TREC topics. To perform parameter training, we randomly divided these documents into two disjoint document pools. We then simulated input from an information retrieval application by generating 15 document sets from each pool, where each set consisted of those documents containing one of 15 arbitrarily chosen keywords; this yielded a total of 30 document sets. Document sets from one pool were used as parameter training data for the sets from the other pool, and vice versa. Performance results are averages over these 30 runs. The scaling factor for IRR was determined by AUTO-SCALE in all cases (again with the same constants  $\alpha$  and  $\beta$  as before). The term-document matrices were created in the same manner as in Section 5.1.




Figure 7: Ten-set averages of kappa average precision and floor and ceiling clustering results, five topics.

Metric	$\kappa$		Clustering floor/ceiling	
	yes	no	yes	no
Is $k$ given?	yes	no	yes	no
Choice of $\ell$	$k$	trained	$k$	trained
# of clusters	N/A		$k$	$\ell$
Section (Data)	§6.1 (Table 9)		§6.2 (Figure 10)	

Figure 8: Evaluation settings for unrestricted distributions. Recall that  $k$  is the number of topics and  $\ell$  is the dimensionality.

## 6.1 Kappa average precision results

Recall that we consider two ways to choose the dimensionality of a document subspace. In the first case, the system knows  $k$ , the number of topics underlying the collection (in practice, this information could be user-supplied as a way to control topic granularity, or given by a set of predetermined classification labels), and sets the dimensionality to it. In the second case,  $k$  is considered unknown, so we simply train the dimensionality parameter using the residual ratio method described in Section 3.3.

From Figure 9, we see that IRR yields higher  $\kappa$  than LSI and VSM for both dimensionality selection methods, and therefore does a better job at representing inter-document similarities. LSI performs relatively poorly on this task; indeed, using  $k$  dimensions in the LSI case leads to worse results than VSM.

Dimensionality selection	LSI	IRR
Number of topics	-8.7	1.4
Trained	0	4.0

Figure 9: Thirty-set average absolute improvement in  $\kappa$  over VSM (51.4%), unrestricted distributions.

## 6.2 Clustering results

To derive floor and ceiling clustering performance results, there are two parameters we need to specify: the dimensionality of the subspace, and the number of clusters.

If  $k$ , the number of topics, is available, then it is the natural choice for the number of clusters. Then, to choose the dimensionality in this case, one option is to also set it to  $k$ ; Figure 10(a) shows the results. We see that IRR has the best clustering performance overall. Note that LSI’s ceiling is actually lower than VSM’s.


When  $k$  is given but we train the dimensionality via the residual ratio, IRR still provides a better subspace for all the clustering algorithms we considered, both in terms of floor and ceiling performance (Figure 10(b)). We observe that for this type of data, training the dimensionality allows LSI to produce improved ceiling results.

We now consider the case in which  $k$  is unknown. In this situation, we know of no alternative but to train the dimensionality on held-out data. As for the number of clusters, a reasonable default is to simply set this value to the trained dimensionality. Of course, this doesn’t apply to VSM, since the dimensionality is not a free parameter for it; instead, we set the number of clusters to the average of the number of topics in the training document sets.

Figure 10(c) shows the clustering results for the unknown- $k$  setting. LSI’s ceiling degrades by 4.3% compared with when the number of topics is given, while those of VSM and IRR show almost no change. Furthermore, IRR clearly outperforms the other methods.

## 6.3 Discussion

In our experiments, LSI did worse or essentially the same as VSM in 4 out of 8 combinations of practical settings and metrics. In particular, when the dimensionality is chosen to be the number of topics, LSI performs relatively poorly. Dimensionality training improves LSI’s kappa average precision scores, and also improves its clustering performance with respect to VSM as long as the correct number of clusters (i.e. the number of topics) is given. However, when



**Figure 10: Document clustering performances, unrestricted distributions: averages over 30 runs.**

the number of clusters is unknown, LSI’s ceiling clustering performance drops, again indicating that for LSI the dimensionality should not be tied to the number of clusters.

In contrast, IRR consistently performs better than LSI and VSM for all our settings and metrics. In particular, IRR fares relatively well when the dimensionality is set to the number of topics as compared to when the dimensionality is actually trained. These results suggest that setting the dimensionality to the number of topics, when known, may be a practical alternative to dimensionality training. Furthermore, in clustering applications for which the number of topics is not known, we at least might be able to reduce the training effort by only searching for the dimensionality, setting the number of clusters to the same value.

## 7. CONCLUSION

To conclude, we review our three main results. First, we have provided a new theoretical analysis of LSI, showing a precise relationship between LSI’s performance and the uniformity of the underlying topic-document distribution. Second, we have used our framework to extend Ando’s (2000) IRR algorithm by giving a novel and effective method for determining the requisite scaling factor. Third, we have shown that IRR, together with our parameter determination method, provides very good performance in comparison to LSI over a variety of document-topic distributions and applications-oriented metrics.

## Acknowledgments

We thank Branimir Boguraev, Roy Byrd, Herb Chong, Jon Kleinberg, Alan Marwick, Mary Neff, John Prager, Edward So, Charlie Van Loan, and Steve Vavasis for many useful discussions, and the anonymous reviewers for their helpful comments. Portions of this work were done while the first author was visiting IBM T. J. Watson Research Center. This paper is based upon work supported in part by the National Science Foundation under ITR/IM grant IIS-0081334. Any opinions, findings, and conclusions or recommendations expressed above are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## 8. REFERENCES

- [1] R. K. Ando. Latent semantic space: Iterative scaling improves inter-document similarity measurement. In *Proceedings of the 23rd SIGIR*, pages 216–223, 2000.
- [2] R. K. Ando. *The Document Representation Problem: An Analysis of LSI and Iterative Residual Rescaling*. PhD thesis, Cornell University, 2001. Forthcoming.
- [3] Y. Azar, A. Fiat, A. Karlin, F. McSherry, and J. Saia. Spectral analysis of data. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2001.
- [4] B. T. Bartell, G. W. Cottrell, and R. K. Belew. Latent Semantic Indexing is an optimal special case of Multidimensional Scaling. In *Proceedings of the 15th SIGIR*, pages 161–167, 1992.
- [5] M. W. Berry, S. T. Dumais, and G. W. O’Brien. Using linear algebra for intelligent information retrieval. *SIAM Review*, 37(4):573–595, 1995.
- [6] B. Boguraev and M. Neff. Discourse segmentation in aid of document summarization. In *Proceedings of the Hawaii International Conference on System Sciences*, 2000.
- [7] J. Carletta. Assessing agreement on classification tasks: The kappa statistic. *Computational Linguistics*, 22(2):249–254, 1996.
- [8] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*, 7(1):1–46, March 1970.
- [9] S. Deerwester, S. Dumais, G. Furnas, T. Landauer, and R. Harshman. Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science*, 41(6):391–407, 1990.
- [10] C. H. Ding. A similarity-based probability model for Latent Semantic Indexing. In *Proceedings of the 22nd SIGIR*, pages 58–65, 1999.
- [11] S. T. Dumais. Improving the retrieval of information from external sources. *Behavior Research Methods, Instruments, & Computers*, 23(2):229–236, 1991.
- [12] S. T. Dumais and J. Nielsen. Automating the assignment of submitted manuscripts to reviewers. In *Proceedings of SIGIR ’92*, pages 233–244, 1992.
- [13] P. W. Foltz and S. T. Dumais. Personalized information delivery: An analysis of information filtering methods. *Communications of the ACM*, 35(12):51–60, 1992.
- [14] P. W. Foltz, W. Kintsch, and T. K. Landauer. The measurement of textual coherence with Latent Semantic Analysis. *Discourse Processes*, 25(2&3):285–307, 1998.
- [15] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, third edition, 1996.
- [16] C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent Semantic Indexing: A probabilistic analysis. *Journal of Computer and System Sciences*, 61(2):217–235, 2000.
- [17] S. Siegel and N. J. Castellan, Jr. *Nonparametric Statistics for the Behavioral Sciences*. McGraw Hill, second edition, 1988.
- [18] N. Slonim and N. Tishby. Document clustering using word clusters via the information bottleneck method. In *Proceedings of the 23rd SIGIR*, pp. 208–215, 2000.



- [19] G. W. Stewart and J. Sun. *Matrix Perturbation Theory*. Computer Science and Scientific Computing. Academic Press, San Diego, 1990.
- [20] R. E. Story. An explanation of the effectiveness of Latent Semantic Indexing by means of a Bayesian regression model. *Information Processing & Management*, 32(3):329–344, 1996.
- [21] M. B. W. Wolfe, M. E. Schreiner, B. Rehder, D. Laham, P. Foltz, W. Kintsch, and T. K. Landauer. Learning from text: Matching readers and text by Latent Semantic Analysis. *Discourse Processes*, 25:309–336, 1998.
- [22] H. Zha, O. Marques, and H. D. Simon. Large-scale SVD and subspace-based methods for information retrieval. In *Solving Irregularly Structured Problems in Parallel, Proceedings of 5th International Symposium (IRREGULAR'98)*, number 1457 in Lecture Notes in Computer Science, pages 29–42. Springer-Verlag, 1998.

## A. APPENDIX

We use the following perturbation result on singular values, which is a slight rewriting of Corollary 8.6.2 in [15].

**THEOREM A.1.** *Let  $\mathbf{E} = \mathbf{X}_1 - \mathbf{X}_2$ , where  $\mathbf{X}_1, \mathbf{X}_2 \in \mathbb{R}^{r \times s}$  and  $r \geq s$ , and for  $1 \leq i \leq s$ , let  $\sigma_i^{(1)}$  and  $\sigma_i^{(2)}$  denote the  $i$ th singular values of  $\mathbf{X}_1$  and  $\mathbf{X}_2$ , respectively. Then,  $|\sigma_i^{(1)} - \sigma_i^{(2)}| \leq \|\mathbf{E}\|_2 \leq \|\mathbf{E}\|_F$ .*

**Notational conventions** If we refer to a singular value  $\sigma_j$  of a matrix  $\mathbf{Z}$  with rank less than  $j$ , it is understood that  $\sigma_j = 0$ . (This corresponds to the “zero-padded” version of the SVD that is also commonly used; we presented the non-padded version above for conceptual clarity.)

Throughout,  $\tau_i$  denotes the  $i$ th singular value of  $\mathbf{P}_{opt}(\mathbf{D})$ .

### A.1 Proof of Theorem 2.1

Let  $\rho_1, \dots, \rho_n$  be the singular values of the true topic-based similarities matrix  $\mathbf{S}$ . For convenience, fix a particular  $i \in \{1, \dots, n\}$ .

Since  $\text{diff}_{\mathbf{S}, \mathbf{D}}(\mathcal{X}_{opt}) = \mathbf{S} - \mathbf{P}_{opt}(\mathbf{D})^T \mathbf{P}_{opt}(\mathbf{D})$ , by Theorem A.1 we have  $\tau_i^2 \in \rho_i \pm \epsilon_{opt}$ .

Next, define the matrix  $\mathbf{S}' \in \mathbb{R}^{n \times n}$  by

$$\mathbf{S}'[t_1, t_2] = \sum_{d \in \mathcal{C}} \text{rel}(t_1, d) \text{rel}(t_2, d),$$

where we define  $\text{rel}(t, d) = 0$  for  $t > k$  and  $d \in \mathcal{C}$ . One can verify that the singular values of  $\mathbf{S}'$  are the same as the singular values of  $\mathbf{S}$ :  $\rho_1, \dots, \rho_n$ .

Now, consider the matrix  $\mathbf{E} \stackrel{\text{def}}{=} \mathbf{S}' - \text{diag}(\Delta_1^2, \dots, \Delta_n^2)$ . We note that the singular values of  $\text{diag}(\Delta_1^2, \dots, \Delta_n^2)$  are  $\Delta_1^2, \dots, \Delta_n^2$ . Furthermore, observe that  $\|\mathbf{E}\|_F = \mu(\mathcal{C})$ , the topic mingling in the collection. Therefore, by again applying Theorem A.1 we find that  $\rho_i \in \Delta_i^2 \pm \mu(\mathcal{C})$ .

Combining our partial results yields the desired result:

$$\tau_i^2 \in \Delta_i^2 \pm (\epsilon_{opt} + \mu(\mathcal{C})).$$

□

### A.2 Invariant subspace tangent theorem

The proof of Theorem 2.2 is based on the following simplification of the *Davis-Kahan tangent theorem* [8] (see also Theorem 3.10 of [19]).

**THEOREM A.2.** *Let  $\mathbf{A} \in \mathbb{R}^{r \times r}$  be a symmetric matrix, and let  $[\mathbf{X} \ \mathbf{Y}]$  be an orthogonal matrix, with  $\mathbf{X} \in \mathbb{R}^{r \times p}$ , so that  $\text{range}(\mathbf{X})$  forms an invariant subspace of  $\mathbf{A}$  (i.e.  $\mathbf{x} \in \text{range}(\mathbf{X})$  implies  $\mathbf{A}\mathbf{x} \in \text{range}(\mathbf{X})$ ). For any matrix  $\hat{\mathbf{X}} \in \mathbb{R}^{r \times p}$  with orthonormal columns, we define the residual matrix  $\mathbf{R}$  of  $\hat{\mathbf{X}}$  as*

$$\mathbf{R} = \mathbf{A}\hat{\mathbf{X}} - (\hat{\mathbf{X}}\hat{\mathbf{X}}^T)\mathbf{A}\hat{\mathbf{X}}.$$

*Suppose the eigenvalues of  $\hat{\mathbf{X}}^T \mathbf{A} \hat{\mathbf{X}}$  lie in the range  $[\alpha, \beta]$  and that there exists  $\delta > 0$  such that the eigenvalues of  $\mathbf{Y}^T \mathbf{A} \mathbf{Y}$  either all lie in the interval  $(-\infty, \alpha - \delta]$  or are all in  $[\beta + \delta, \infty)$ . Then,*

$$\|\tan(\Theta(\text{range}(\mathbf{X}), \text{range}(\hat{\mathbf{X}}))\|_2 \leq \frac{\|\mathbf{R}\|_2}{\delta}.$$

### A.3 Proof sketch for Theorem 2.2

Here, we outline the proof (given in full in [2]). The main idea is to apply Theorem A.2 by choosing  $\mathbf{X}$  and  $\hat{\mathbf{X}}$  so that  $\text{range}(\mathbf{X}) = \mathcal{X}_{LSI}$  and  $\text{range}(\hat{\mathbf{X}}) = \mathcal{X}_{opt}$ , and setting  $\mathbf{A} = \mathbf{D}\mathbf{D}^T$ , for which the LSI subspace is invariant.

Let  $\hat{\mathbf{D}} = \mathbf{P}_{opt}(\mathbf{D})$ , and define  $\bar{\mathbf{D}} = \mathbf{D} - \hat{\mathbf{D}}$ ; note that  $\hat{\mathbf{D}}^T \bar{\mathbf{D}} = \bar{\mathbf{D}}^T \hat{\mathbf{D}} = \mathbf{0}$ . Let  $\sigma_i$  denote the  $i$ th singular value of  $\mathbf{D}$ .

First, consider the largest singular value of  $\bar{\mathbf{D}}^T \bar{\mathbf{D}}$ , which we will denote by  $\hat{\epsilon}_{vsm}$ . It can be shown that  $\hat{\epsilon}_{vsm} \in \epsilon_{vsm} \pm \epsilon_{opt}$ , thus justifying our choice of notation. (First show that  $\text{diff}_{\mathbf{D}}(\mathcal{X}_{opt}) = \text{diff}_{\mathbf{S}, \mathbf{D}}(\mathcal{X}_{VSM}) - (-\bar{\mathbf{D}}^T \bar{\mathbf{D}})$ , and then apply Theorem A.1 to this equation and observe that the largest singular value of  $\text{diff}_{\mathbf{S}, \mathbf{D}}(\mathcal{X}_{VSM})$  is  $\epsilon_{vsm}$ .)

Then, it can be shown that  $\sigma_{h+1} \leq \sqrt{\hat{\epsilon}_{vsm}}$ : apply Theorem A.1 to the equation  $\bar{\mathbf{D}}^T \bar{\mathbf{D}} = \mathbf{D}^T \mathbf{D} - \hat{\mathbf{D}}^T \hat{\mathbf{D}}$  and then note that  $\tau_{h+1} = 0$  because of the dimensionality of  $\mathcal{X}_{opt}$ .

Now, to apply Theorem A.2, choose the columns of  $\mathbf{X}$  to be the first  $h$  left singular vectors (in order) of  $\mathbf{D}$ , and choose the columns of  $\hat{\mathbf{X}}$  to be all  $h$  left singular vectors (in order) of  $\hat{\mathbf{D}}$ . We set the columns of  $\mathbf{Y}$  to the rest of the  $m - h$  left singular vectors (in order) of the “zero-padded” version of the SVD of  $\mathbf{D}$ .

Some linear algebra reveals that the eigenvalues of  $\mathbf{Y}^T \mathbf{A} \mathbf{Y}$  are no greater than  $\sigma_{h+1}^2$ , and the eigenvalues of  $\hat{\mathbf{X}}^T \mathbf{A} \hat{\mathbf{X}}$  are no smaller than  $\hat{\Delta}_{min}^2$ . Therefore, we set  $\delta = \hat{\Delta}_{min}^2 - \sigma_{h+1}^2$ . Note that  $\delta \geq \hat{\Delta}_{min}^2 - \hat{\epsilon}_{vsm}$  by above, and so is positive by assumption. Hence, Theorem A.2 applies.

Finally, it can be shown that  $\|\mathbf{R}\|_2 \leq \hat{\Delta}_{max} \cdot \sqrt{\hat{\epsilon}_{vsm}}$ , which yields the desired result:

$$\begin{aligned} \|\tan(\Theta(\mathcal{X}_{LSI}, \mathcal{X}_{opt}))\|_2 &\leq \frac{\hat{\Delta}_{max} \cdot \sqrt{\hat{\epsilon}_{vsm}}}{\hat{\Delta}_{min}^2 - \hat{\epsilon}_{vsm}} \\ &\leq \frac{\hat{\Delta}_{max}}{\hat{\Delta}_{min}} \cdot \frac{\sqrt{\hat{\epsilon}_{vsm}} / \hat{\Delta}_{min}}{1 - (\sqrt{\hat{\epsilon}_{vsm}} / \hat{\Delta}_{min})^2}. \end{aligned}$$

□