

(c) If $C = \begin{pmatrix} 1 & 2 \\ 2 & 3 \end{pmatrix}$, then $C^T = \begin{pmatrix} 1 & 2 \\ 2 & 3 \end{pmatrix}$. ■

The matrix C in Example 11 is its own transpose. This frequently happens with matrices that arise in applications.

Definition

An $n \times n$ matrix A is said to be **symmetric** if $A^T = A$.

The following are some examples of symmetric matrices:

$$\begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix} \quad \begin{pmatrix} 2 & 3 & 4 \\ 3 & 1 & 5 \\ 4 & 5 & 3 \end{pmatrix} \quad \begin{pmatrix} 0 & 1 & 2 \\ 1 & 1 & -2 \\ 2 & -2 & -3 \end{pmatrix}$$

APPLICATION 2 Information Retrieval

The growth of digital libraries on the Internet has led to dramatic improvements in the storage and retrieval of information. Modern retrieval methods are based on matrix theory and linear algebra.

In a typical situation, a database consists of a collection of documents and we wish to search the collection and find the documents that best match some particular search conditions. Depending on the type of database, we could search for such items as research articles in journals, Web pages on the Internet, books in a library, or movies in a film collection.

To see how the searches are done, let us assume that our database consists of m documents and that there are n dictionary words that can be used as keywords for searches. Not all words are allowable, since it would not be practical to search for common words such as articles or prepositions. If the key dictionary words are ordered alphabetically, then we can represent the database by an $m \times n$ matrix A . Each document is represented by a column of the matrix. The first entry in the j th column of A would be a number representing the relative frequency of the first key dictionary word in the j th document. The entry a_{2j} represents the relative frequency of the second word in the j th document, and so on. The list of keywords to be used in the search is represented by a vector \mathbf{x} in \mathbb{R}^m . The i th entry of \mathbf{x} is taken to be 1 if the i th word in the list of keywords is on our search list; otherwise, we set $x_i = 0$. To carry out the search, we simply multiply A^T times \mathbf{x} .

Simple Matching Searches

The simplest type of search determines how many of the key search words are in each document; it does not take into account the relative frequencies of the words. Suppose, for example, that our database consists of these book titles:

- B1.** *Applied Linear Algebra*
- B2.** *Elementary Linear Algebra*
- B3.** *Elementary Linear Algebra with Applications*
- B4.** *Linear Algebra and Its Applications*
- B5.** *Linear Algebra with Applications*

Year
8,060
14,780
7,140
29,980

and addition
lications are
lications.

ose columns

(8)

tively, as the
tively, as the

B6. Matrix Algebra with Applications

B7. Matrix Theory

The collection of keywords is given by the following alphabetical list:

algebra, application, elementary, linear, matrix, theory

For a simple matching search, we just use 0's and 1's, rather than relative frequencies for the entries of the database matrix. Thus, the (i, j) entry of the matrix will be 1 if the i th word appears in the title of the j th book and 0 if it does not. We will assume that our search engine is sophisticated enough to equate various forms of a word. So, for example, in our list of titles the words *applied* and *applications* are both counted as forms of the word *application*. The database matrix for our list of books is the array defined by Table 4.

Table 4 Array Representation for Database of Linear Algebra Books

Key Words	Books						
	B1	B2	B3	B4	B5	B6	B7
<i>algebra</i>	1	1	1	1	1	1	0
<i>application</i>	1	0	1	1	1	1	0
<i>elementary</i>	0	1	1	0	0	0	0
<i>linear</i>	1	1	1	1	1	0	0
<i>matrix</i>	0	0	0	0	0	1	1
<i>theory</i>	0	0	0	0	0	0	1

If the words we are searching for are *applied*, *linear*, and *algebra*, then the database matrix and search vector are respectively given by

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad \mathbf{x} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

If we set $\mathbf{y} = A^T \mathbf{x}$, then

$$\mathbf{y} = \begin{pmatrix} 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 3 \\ 2 \\ 3 \\ 3 \\ 3 \\ 2 \\ 0 \end{pmatrix}$$

The value of y_1 is the number of search word matches in the title of the first book, the value of y_2 is the number of matches in the second book title, and so on. Since

$y_1 = y_3 = y_4 = y_5 = 3$, the titles of books B1, B3, B4, and B5 must contain all three search words. If the search is set up to find titles matching all search words, then the search engine will report the titles of the first, third, fourth, and fifth books.

Relative-Frequency Searches

Searches of noncommercial databases generally find all documents containing the key search words and then order the documents based on the relative frequencies of the keywords. In this case, the entries of the database matrix should represent the relative frequencies of the keywords in the documents. For example, suppose that in the dictionary of all key words of the database the 6th word is *algebra* and the 8th word is *applied*, where all words are listed alphabetically. If, say, document 9 in the database contains a total of 200 occurrences of keywords from the dictionary, and if the word *algebra* occurred 10 times in the document and the word *applied* occurred 6 times, then the relative frequencies for these words would be $\frac{10}{200}$ and $\frac{6}{200}$, and the corresponding entries in the database matrix would be

$$a_{69} = 0.05 \quad \text{and} \quad a_{89} = 0.03$$

To search for these two words, we take our search vector \mathbf{x} to be the vector whose entries x_6 and x_8 are both equal to 1 and whose remaining entries are all 0. We then compute

$$\mathbf{y} = A^T \mathbf{x}$$

The entry of \mathbf{y} corresponding to document 9 is

$$y_9 = a_{69} \cdot 1 + a_{89} \cdot 1 = 0.08$$

Note that 16 of the 200 words (8% of the words) in document 9 match the key search words. If y_j is the largest entry of \mathbf{y} , this would indicate that the j th document in the database is the one that contains the keywords with the greatest relative frequencies.

Advanced Search Methods

A search for the keywords *linear* and *algebra* could easily turn up hundreds of documents, some of which may not even be about linear algebra. If we were to increase the number of search words and require that all search words be matched, then we would run the risk of excluding some crucial linear algebra documents. Rather than match all words of the expanded search list, our database search should give priority to those documents which match most of the keywords with high relative frequencies. To accomplish this, we need to find the columns of the database matrix A that are “closest” to the search vector \mathbf{x} . One way to measure how close two vectors are is to define *the angle between the vectors*. We will do this in Section 1 of Chapter 5.

We will also revisit the information retrieval application after we have learned about the *singular value decomposition* (Chapter 6, Section 5). This decomposition can be used to find a simpler approximation to the database matrix, which will speed up the searches dramatically. Often it has the added advantage of filtering out *noise*; that is, using the approximate version of the database matrix may automatically have the effect of eliminating documents that use keywords in unwanted contexts. For example, a dental student and a mathematics student could both use *calculus* as one of their

search words. Since the list of mathematics search words does not contain any other dental terms, a mathematics search using an approximate database matrix is likely to eliminate all documents relating to dentistry. Similarly, the mathematics documents would be filtered out in the dental student's search.

Web Searches and Page Ranking

Modern Web searches could easily involve billions of documents with hundreds of thousands of keywords. Indeed, as of July 2008, there were more than 1 trillion Web pages on the Internet, and it is not uncommon for search engines to acquire or update as many as 10 million Web pages in a single day. Although the database matrix for pages on the Internet is extremely large, searches can be simplified dramatically, since the matrices and search vectors are *sparse*; that is, most of the entries in any column are 0's.

For Internet searches, the better search engines will do simple matching searches to find all pages matching the keywords, but they will not order them on the basis of the relative frequencies of the keywords. Because of the commercial nature of the Internet, people who want to sell products may deliberately make repeated use of keywords to ensure that their Web site is highly ranked in any relative-frequency search. In fact, it is easy to surreptitiously list a keyword hundreds of times. If the font color of the word matches the background color of the page, then the viewer will not be aware that the word is listed repeatedly.

For Web searches, a more sophisticated algorithm is necessary for ranking the pages that contain all of the key search words. In Chapter 6, we will study a special type of matrix model for assigning probabilities in certain random processes. This type of model is referred to as a *Markov process* or a *Markov chain*. In Section 3 of Chapter 6, we will see how to use Markov chains to model Web surfing and obtain rankings of Web pages.

References

1. Berry, Michael W., and Murray Browne, *Understanding Search Engines: Mathematical Modeling and Text Retrieval*, SIAM, Philadelphia, 1999.

SECTION 1.3 EXERCISES

1. If

$$A = \begin{bmatrix} 3 & 1 & 4 \\ -2 & 0 & 1 \\ 1 & 2 & 2 \end{bmatrix} \text{ and } B = \begin{bmatrix} 1 & 0 & 2 \\ -3 & 1 & 1 \\ 2 & -4 & 1 \end{bmatrix}$$

compute

- | | |
|---------------|-----------------------|
| (a) $2A$ | (b) $A + B$ |
| (c) $2A - 3B$ | (d) $(2A)^T - (3B)^T$ |
| (e) AB | (f) BA |
| (g) $A^T B^T$ | (h) $(BA)^T$ |

2. For each of the pairs of matrices that follow, determine whether it is possible to multiply the first matrix times the second. If it is possible, perform the multiplication.

(a) $\begin{bmatrix} 3 & 5 & 1 \\ -2 & 0 & 2 \end{bmatrix} \begin{bmatrix} 2 & 1 \\ 1 & 3 \\ 4 & 1 \end{bmatrix}$

(b) $\begin{bmatrix} 4 & -2 \\ 6 & -4 \\ 8 & -6 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \end{bmatrix}$

(c) $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$

(d) $\begin{bmatrix} 4 \\ 2 \end{bmatrix}$

(e) $\begin{bmatrix} 4 \\ 2 \end{bmatrix}$

(f) $\begin{bmatrix} 2 \\ -1 \\ 3 \end{bmatrix}$

3. For which multiply t would the

4. Write eac a matrix e

(a) $3x_1 + 2x_1 -$

(b) $x_1 + 2x_1 + 3x_1 -$

(c) $2x_1 + x_1 - 3x_1 -$

5. If

verify that

(a) $5A =$

(c) $(A^T)^T$

6. If

$$A = \begin{bmatrix} 4 \\ 2 \end{bmatrix}$$

verify that

(a) $A + B$

(b) $3(A +$

(c) $(A + I$

7. If

$$A = \begin{bmatrix} \\ \\ \end{bmatrix}$$

verify that

(a) $3(AB)$

(b) $(AB)^T$