

CLAIMS

What is claimed is:

1. A computerized method of compressing symbolic information organized into a plurality of documents, each document having a plurality of symbols, the method comprising:
 - (a) identifying a plurality of symbol pairs, each symbol pair consisting of two sequential or non-sequential symbols in the input document, one or more symbol pairs consisting of two non-sequential symbols in the input document;
 - (b) for each unique symbol pair of the plurality of symbol pairs, updating a count identifying the number of appearances of the unique symbol pair; and
 - (c) producing a compressed document by causing the compressed document to include, at each position associated with one of the plurality of symbol pairs from the input document, including one or more symbol pairs consisting of two non-sequential symbols, (i) a replacement symbol associated by a compression dictionary with the unique symbol pair matching the one of the plurality of symbol pairs, if the count for the unique symbol pair exceeds a threshold, and (ii) for at least those symbol pairs consisting of two non-sequential symbols, indicia indicating a distance between locations of the non-sequential symbols of the pair in the input document.
2. The method of claim 1 wherein the plurality of symbol pairs includes only those pairs of non-sequential symbols for which the distance between locations of the non-sequential symbols of the pair in the input document is less than a numeric distance cap.
3. The method of claim 2 wherein the numeric distance cap is ten symbols.
4. The method of claim 1 wherein the threshold is two.
5. The method of claim 1 further comprising sorting the compression dictionary in order of the count for the respective unique symbol pairs.
6. The method of claim 5 further comprising excluding or removing from the compression dictionary all unique symbol pairs having a count that does not exceed the threshold and making an association in the compression dictionary

between each unique symbol pair having a count that exceeds the threshold and a replacement symbol.

7. The method of claim 1 further comprising performing an additional pass by repeating parts (a), (b), and (c) using as the input document the compressed document.

8. The method of claim 7 further comprising performing additional passes, until the compressed document created by one of the passes is unaltered from the compressed document created by the immediately previous pass.

9. The method of claim 1 wherein each symbol from the first document indicates a word.

10. The method of claim 9 further comprising, before part (a), creating the first document by replacing each word of a text document with a corresponding one of the plurality of symbols.

11. The method of claim 1 wherein each symbol in the first document indicates an alphanumeric character.

12. The method of claim 1 wherein each replacement symbol comprises an address of an entry in the compression dictionary, which entry stores the unique symbol pair associated with replacement symbol.

13. The method of claim 1 further comprising performing parts (a), (b), and (c) on each of the other documents of the plurality of documents, wherein the respective counts for the symbol pairs identifies the number of previous appearances of that symbol pair in any of the plurality of documents.

14. The method of claim 13 wherein parts (a) and (b) are performed on each of the plurality of documents before part (c) is performed on any of the plurality of documents.

15. The method of claim 14 further comprising, before part (c) is performed on any of the plurality of documents, making an association in the compression dictionary between each unique symbol pair and a replacement symbol if and only if the unique symbol pair has a count that exceeds the threshold.

16. The method of claim 13 further comprising sorting the compression dictionary in order of the count for the respective unique symbol pairs.

17. The method of claim 16 further comprising removing from the compression dictionary all unique symbol pairs having a count that does not exceed the threshold and making an association in the compression dictionary

between each unique symbol pair having a count that exceeds the threshold and a replacement symbol.

18. The method of claim 13 further comprising repeating parts (a), (b), and (c) on each of the other documents to make at least one additional pass by
5 repeating parts (a), (b), and (c), using as the input document the compressed document previously produced by the process, until the compressed document created by one of the passes is unaltered from the compressed document created by the immediately previous pass.

19. The method of claim 18 further comprising assigning the first document
10 at least one of a plurality of attributes based on the frequency with which a replacement symbol associated with a unique symbol pair appears in an output document associated with the first document.

20. The method of claim 19 further comprising suggesting a related document based the at least one attribute of the first document and the at least one
15 attribute of the related document.

21. A computer system having at least one processor capable of compressing symbolic information organized into a plurality of documents, each document having a plurality of symbols, the computer system programmed to perform automatically, with the processor, the method of any one of claims 1
20 through 20 with a first document of the plurality of documents as an input document.

22. A computer-readable storage medium that is not a transitory propagating signal storing a set of computer instructions for compressing symbolic information organized into a plurality of documents, each document having a
25 plurality of symbols, wherein the set of computer instructions, when executed on the computer, causes the computer to perform automatically the method of any one of claims 1 through 20 with a first document of the plurality of documents as an input document.