# Word Hy-phen-a-tion by Com-put-er 

 byFranklin Mark Liang

## Department of Computer Science

Stanford University
Stanford, CA 94305

# WORD HY-PHEN-A-TION BY COM-PUT-ER 

Franklin Mark Liang<br>Department of Computer Science<br>Stanford University<br>Stanford, California 94305

Abstract
This thesis describes research leading to an improved word hyphenation algorithm for the $T_{\mathbb{E}} \mathrm{X} 82$ typesetting system. Hyphenation is viewed primarily as a data compression problem, where we are given a dictionary of words with allowable division points, and try to devise methods that take advantage of the large amount of redundancy present.

The new hyphenation algorithm is based on the idea of hyphenating and inhibiting patterns. These are simply strings of letters that, when they match in a word, give us information about hyphenation at some point in the pattern. For example, '-tion' and ' $c-c$ ' are good hyphenating patterns. An important feature of this method is that a suitable set of patterns can be extracted automatica!!y from the dictionary.

In order to represent the set of patterns in a compart form that is also reasonably efficient for searching, the author has developed a new data structure called a packed trie. This data structure allows the very fast search times characteristic of indexed tries, but in many cases it entirely eliminates the wasted space for null links usually present in such tries. We demonstrate the versatility and practical advantages of this data structure by using a variant of it as the critical component of the program that generates the patterns from the dictionary.

The resulting hyphenation algorithm uses about 4500 patterns that compile into a packed trie occupying 25 K bytes of storage. These patterns find $89 \%$ of the hyphens in a pocket dictionary word list, with essentially no error. By comparison, the uncompressed dictionary occupies over 500 K bytes.

This research was supported in part by the National Science Foundation under grants IST-8801926 and MSC-89-00984, and by the System Development Foundation. ' $T_{E} X$ ' is a trademark of the American Mathematical Society.

# WORD HY-PHEN-A-TION BY COM-PUT-ER 

A DISSERTATION<br>SUBMITTED TO THE DRPARTMENT OF COMPUTER SCIENCE AND THE COMMITTEE ON GRADUATE STUDIES<br>OF STANFORD UNIVERSITY<br>ES PARTIAL FULFILLMENT OF THE REQUIREMENTS<br>FOR THE DEGREE OF<br>DOCTOR OF PHILOSOPHY

by<br>Franklin Mark Liang<br>June 1983

(c) Copyright 1983

## by

Franklin Mark Liang

## Acknowledgments

I am greatly indebted to my adviser, Donald Knuth, for creating the research environment that made this work possible. When I began work on the $\mathrm{T}_{\mathrm{E}} \mathrm{X}$ project as a summer job, I would not have predicted that computer typesetting would become such an active area of computer science research. Prof. Knuth's foresight was to recognize that there were a number of fascinating problems in the field waiting to be explored, and his pioneering efforts have stimulated many others to think about these problems.

I am also grateful to the Stanford Computer Science Department for providing the facilities and the community that have formed the major part of my life for the past several years.

I thank my readers, Luis Trabb Pardo and John Gill, as well as Leo Guibas who served on my orals committee on short notice.

In addition, thanks to David Fuchs and Tom Pressburger for helpful advice and encouragement.

Finally, this thesis is dedicated to my parents, for whom the experience of pursuing a graduate degree has been perhaps even more traumatic than it was for myself.

## Table of contents

Introduction ..... 1
Examples ..... 2
$\mathrm{T}_{\mathrm{E}} \mathrm{X}$ and hyphenation ..... 3
Time magazine algorithm ..... 4
Patterns ..... 5
Overview of thesis ..... 7
The dictionary problem ..... 8
Data structures ..... 9
Superimposed coding ..... 10
Tries ..... 11
Packed tries ..... 15
Suffix compression ..... 16
Derived forms ..... 18
Spelling checkers ..... 19
Related work ..... 21
Hyphenation ..... 23
Finite-state machines with output ..... 23
Minimization with don't cares ..... 24
Pattern matching ..... 26
Pattern generation ..... 29
Heuristics ..... 30
Collecting pattern statistics ..... 31
Dynamic packed tries ..... 32
Experimental results ..... 34
Examples ..... 37
History and Conclusion ..... 39
Appendix ..... 45
The PATGEN program ..... 45
List of patterns ..... 74
Reforences ..... 83

## Chapter 1

## Introduction

[^0]amenable to computer implementation;.e.g. "split between the elements of a compound word". Rule-based schemes are inevitably subject to error, and they rarely cover all possible cases. In addition, the task of finding a suitable set of rules in the first place can be a difficult and lengthy project.

Dictionary-based routines simiply store an entire word list along with the allowable division points. The obvious disadvantage of this method is the excessive storage required, as well as the slowing down of the justification process when the hyphenation routine needs to access a part of the dictionary on secondary store.

## Examples

To demonstrate the importance of hyphenation, consider Figure 1, which shows a paragraph set in three different ways by $\mathrm{T}_{\mathrm{E}} \mathrm{X}$. The first example uses $\mathrm{T}_{\mathrm{E}} \mathrm{X}$ 's normal paragraph justification parameters, but with the hyphenation routine turned off. Because the line width in this example is rather narrow, $\mathrm{T}_{\mathrm{E}} \mathrm{X}$ is unable to find an acceptable way of justifying the paragraph, resulting in the phenomenon known as an "overfull box".

One way to fix this problem is to increase the "stretchability" of the spaces between words, as shown in the second example. (TEX users: This was done by increasing the stretch component of spaceskip to .5 em .) The right margin is now straight, as desired, but the overall spacing is somewhat loose.

In the third example, the hyphenation routine is turned on, and everything is beautiful.


In olden times when wishing still helped one, there lived a king whose daughters were all beautiful, but the youngest was so benutiful that the sun itself, which has seen so much, was astonished whenever it shone in her face. Close by the king's castle lay a great dark fcrest, and under $a$ old lime-tree in the forest was a well, and when the day was very warm, the king's child went out into the forest and snt down by the side of the cool fountain, and when she was bored she took a golden ball, and threw it up on high and caught it, and this ball was her favorite plaything.

In olden times when wishing atill helped one, there lived a king whose daughters were all beautiful, but the youngest was so beautiful that the sun itself, which has seen so much, was astonished whenever it shone in her face. Close by the king's castle lay a great dark forest, and under an old lime-tree in the forest was a well, and when the day was very warm, the king's child went out into the forest and sat down by the side of the cool fountain, and when she was bored she took a golden ball, and threw it up on high and cnupht it, and this ball was her favorite plaything.

Figure 1. A typical paragraph with and without hyphenation.

| sel-fadjoint | as-so-ciate | as-so-ci-ate |
| :--- | :--- | :--- |
| Pit-tsburgh | prog-ress | pro-gress |
| clearin-ghouse | rec-ord | re-cord |
| fun-draising | a-rith-me-tic | ar-ith-met-ic |
| ho-meowners | eve-ning | even-ing |
| playw-right | pe-ri-od-ic | per-i-o-dic |
| algori-thm |  |  |
| walkth-rough | in-de-pen-dent | in-de-rend-ent |
| Re-agan | tri-bune | trib-une |

Figure 2. Difficult hyphenations.
However, life is not always so simple. Figure 2 shows that hyphenation can be difficult. The first column shows erroneous hyphenations made by various typesetting systems (which shall remain nameless). The next group of examples are words that hyphenate differently depending on how they are used. This happens most commonly with words that can serve as both nouns and verbs. The last two examples show that different dictionarics do not always agree on hyphenaiion (in this case Webster's vs. American Heritage).
$\mathrm{T}_{\mathrm{E}} \mathrm{X}$ and hyphenation
The original $T_{E} X$ hyphenation algorithm was designed by Prof. Knuth and the author in the summer of 1977. It is essentially a rule-based algorithm, with three main types of rules: (1) suffix removal, (2) prefix removal, and (3) vowel-consonant-consonant-vowel (vccv) breaking. The latter rule states that when the pattern 'vowel-consonant-consonant-vowel' appears in a word, we can in most cases split between the consonants. There are also many special case rules; for example, "break vowel-q" or "break after ck". Finally a small exception dictionary (about 300 words) is used to handle particularly objectionable errors made by the above rules, and to hyphenate certain common words (e.g. pro-gram) that are not split by the rules. The complete algorithm is described in Appendix H of the old TEX manual.

In practice, the above algorithm has served quite well. Although it does not find all possible division points in a word, it very rarely makes an error. Tests on a pocket dictionary word list indicate that about $40 \%$ of the allowable hyphen points are found, with $1 \%$ error (relative to the total number of hyphen points). The algorithm requires 4 K 36 -bit words of code, including the exception dictionary.

The goal of the present research was to develop a better hyphenation algorithm. By "better" we mean finding more hyphens, with little or no error, and using as little additional space as possible. Recall that one way to perform hyphenation is to simply store the entire dictionary. Thus we can view our task as a data compression problem. Since there is a good deal of redundancy in English, we can hope for substantial improvement over the straightforward representation.

Another goal was to automate the design of the algorithm as much as possible. The original $T_{E} \mathrm{X}$ algorithm was developed mostly by hand, with a good deal of trial and error. Extending such a rule-based scheme to find the remaining hyphens seems very difficult. Furthermore such an effort must be repeated for each new language. The former approach can be a problem even for English, because pronunciation (and thus hyphenation) tends to change over time, and because different types of publication may call for different sets of admissible hyphens.

## Time magazine algorithm

A number of approaches were considered, including methods that have been discussed in the literature or implemented in existing typesetting systems. One of the methods studied was the so-called Time magazine algorithm, which is table-based rather than rule-based.

The idea is to look at four letters surrounding each possible 'reakpoint, namely two letters preceding and two letters following the given point. However we do not want to store a table of $26^{4}=456,976$ entrics representing all possible four-letter combinations. (In practice only about $15 \%$ of these four-letter combinations actually occur in Finglish words, but it is not immediately obvious how to take advantage of this.)

Instead, the method uses three tables of size $26^{2}$, corresponding to the two letters preceding, surrounding, and following a potential hyphen point. That is, if the letter pattern wx-yz occurs in a word, we look up three values corresponding to the letter pairs $w x, x y$, and $y z$, and use these values to determine if we can split the pattern.

What should the three tables contain? In the Tine algorithm the table values were the probabilities that a hyphen could occur after, between, or before two given letters, respectively. The probability that the pattern $w x-y z$ can be split is then estimated as the product of these three values (as if the probabilities were independent, which they aren't). Finally the estimated value is compared against a threshold to determine hyphenation. Figure 3 shows an example of hyphenation probabilities computed by this method.


SUPERCALIFRAGIIISTICEXPIALIDOCIOUS

Figure 3. Hyphenation probabilities.
The advantage of this table-based approach is that the tables can be generated automatically from the dictionary. However, some experiments with the method yielded discouraging results. One estimate is $40 \%$ of the hyphens found, with $8 \%$ error. Thus a large exception dictionary would be required for good performance.

The reason for the limited performance of the above scheme is that just four letters of context surrounding the potential break point are not enough in many cases. In an extreme example, we might have to look as many as 10 letters ahead in order to determine hyphenation, e.g. dem-on-stra-tion vs. de-mon-stra-tive.

So a more powerful method is needed.

## Patterns

A good deal of experimentation led the author to a more powerful method based on the idea of hyphenation patterns. These are simply strings of letters that, when they match in a word, will tell us how to hyphenate at some point in the pattern. For example, the pattern 'tion' might tell us that we can hyphenate before the ' $t$ '. Or when the pattern ' $c c$ ' appears in a word, we can usually hyphenate between the $c$ 's. Here are some more examples of good hyphenating patterns:

```
.in-d .in-s .in-t .un-d b-s -cia con-s con-t e-ly er-1 er-m
ex- -ful it-t i-ty -less l-ly -ment n-co -ness n-f n-1 n-si
n-v om-m -sion s-ly s-nos ti-ca x-p
```

(The character '.' matches the beginning or end of a word.)

Patterns have many advantages. They are a general form of "hyphenation rule" that can include prefix, suffix, and other rules as special cases. Patterns can even describe an exception dictionary, namely by using entire words as patterns. (Actually, patterns are often more concise than an exception dictionary because a single pattern can handle several variant forms of a word; e.g. pro-gram, pro-grams, and pro-grammed.)

More importantly, the pattern matching approach has proven very effective. An appropriate set of patterns captures very concisely the information needed to perform hyphenation. Yet the pattern rules are of simple enough form that they can be generated automatically from the dictionary.

When looking for good hyphenating patterns, we soon discover that almost all of them have some exceptions. Although -tion is a very "safe" pattern, it fails on the word cat-ion. Most other cases are less clear-cut; for example, the common paittern $n-t$ can be hyphenated about \&of percent of the time. It definitely seems worthwhile to use such patterns, provided that we can deal with the exceptions in some manner.

After choosing a set of hyphenating patterns, we may end up with thousands of exceptions. There could be listed in an exception dictionary, but we soon notice there are many similarities among the exceptions. For example, in the original $T_{E} \mathrm{X}$ algorithm we found that the vowel-consonant-consonant-vowel rule resulted in hundreds of errors of the form $X$-Yer or $X$-Yers, for certain consonant pairs $X Y$, so we put in a new rule to prevent those errors.

Thus, there may be "rules" that can handle large classes of exceptions. To take advariage of this, patterns come to the rescue again; but this time they are inhibit-* ir. patterns, because they show where hyphens should not be placed. Some good examples of inhibiting patterns are: $b=1 y$ (don't break between $b$ and $l y$ ), $b s=,=c i n g$, $\mathrm{i}=\mathrm{n}, \mathrm{i}=\mathrm{tin},=1 \mathrm{~s}, \mathrm{nn}=, \mathrm{ns}=\mathrm{t}, \mathrm{n}=\mathrm{ted},=\mathrm{pt}, \mathrm{ti}=\mathrm{al},=\mathrm{tly},=\mathrm{ts}$, and $\mathrm{tt}=$.

As it turns out, this approach is worth pursuing further. That is, after app'ying hyphenating and inhibiting patterns as discussed above, we might have another set of hyphenaing patterns, then another set of inhibiting patterns, and so on. We can think of each level of patterns as being "exceptions to the exceptions" of the previous level. The current $\mathrm{T}_{\mathrm{E}} \mathrm{X} 82$ algorithm uses five alternating levels of hyphenating and inhibiting patterns. The reasons for this will be explained in Chapter 4.

The idea of patterns is the basis of the new TEX hyphenation algorithm, and it was the inspiration for much of the intermediate investigation, that will be described.

## Overview of thesis

In developing the pattern scheme, two main questions arose: (1) How can we represent the set of hyphenation patterns in a compact form that is also reasonably efficient for searching? (2) Given a hyphenated word list, how can we generate a suitable set of patterns?

To solve these problems, the author has developed a new data structure called a parked trie. This data structure alows the very fast search times characteristic of indexed tries, but in many cases it entirely eliminates the wasted space for null links usually present in such tries.

We will demo:astrate the versatility and practical advantages of this data structurs 'y waing it not only to represent the hyphenation patterns in the final algorithm, but also as the critical component of the program that generates the patterns from the distionary. Packed tries have many other potential applications, including identifier lookup, spelling checking, and lexicographic sorting.

Chapter 2 considers the simpler problem of recognizing, rather than hyphenating, a set of words such as a dictionary, and uses this problem to motivate and explain the advantages of the packed trie data structure. We also point out the close relationship between tries and finite-state machines.

Chapter 3 discusses ways of applying these ideas to hyphenation. After considering various approaches, including minimization with don't cares, we return to the idea of patterns.

Chapter 4 discusses the heuristic method used to select patterns, introduces dynamic packed tries, and describes some experiments with the pattern generation pro-gram.

Chapter 5 gives a brief history, and mentions ideas for future research.
Finally, the appendix contains the WEB [3] listing of the portable pattern generation program PATGEN, as well as the set of patterns currently used by TEX82.

Note: The present chapter has been typeset by giving unusual instructions to $\mathrm{T}_{\mathrm{E}} \mathrm{X}$ so that it hyphenates words much more often than usual; therefore the reader can see numerous examples of word breaks that were discovered by the new algorithm.

## Chapter 2

## The dictionary problem

In this chapter we consider the problem of recognizing a set of words over and alphabet. To be more precise, an alphabet is a set of characters or symbols, for example the lucters A through Z , or the ASCII character set. A word is a sequence of characters from the alphabet. Given a set of words, our problem is to design a data structure that will allow us to determine efficiently whether or not some word is in the set.

In particular, we will use spelling checking as an example throughout this chapter. This is a topic of interest in its own right, but we discuss it here because the pattern matching techniques we propose will turn out to be very useful in our hyphenation algorithm.

Our problem is a special case of the general set recognition problem, because the elements of our set have the additional structure of being variable-length sequences of symbols from a finite alphabet. This naturally suggests methods based on a character-by-character examination of the key, rather than methods that operate on the entire key at once. Also, the redundancy present in natural languages such as English suggests additional opportunities for compression of the set representation.

We will be especially interested in space minimization. Most data structures for set representation, including the one we propose, are reasonably fast for searching. That is, a search for a key doesn't take much more time than is needed to examine the key itself. However, most of these algorithms assume that everything is "in core", that is, in the primary memory of the computer. In many situations, such as our spelling checking example, this is not feasible. Since secondary memory access times are typically much longer, it is worthwhile to try compressing the data structure as much as possible.

In addition to determining whether a given word is in the set, there are other operations we might wish to perform on the set representation. The most basic are insertion and deletion of words from the set. More complicated operations include performing the union of two sets, partitioning a set according to some criterion,
determining which of several sets an element is a member of, or operations based on an ordering or other auxiliary information associated with the keys in the set. For the data structures we consider, we will pay some attention to methods for insertion and deletion, but we shall not discuss the more complicated operations.

We first survey some known methods for set representation, and then propose a new data structure called a "packed trie".

## Data structures

Methods for set representation include the following: sequential lists, sorted listz, binary search trees, balanced trees, hashing, superimposed coding, bit vectors, and digital search trees (also known as tries). Good discussions of these data structures can be found in a number of texts, including Knuth [4], Standish [5], and AHU [6]. Below we make a few remarks about each of these representations.

A sequential list is the most straightforward representation. It requires both space and search time proportional to the number of characters in the dictionary.

A sorted list assumes an ordering on the keys, such as alphabetical order. Binary search allows the search time to be reduced to the logarithm of the size of the dictionary, but space is not reduced.

A binary search tree also allows search in logarithmic time. This can be thought of as a more flexible version of a sorted list that can be optimized in various ways. For example if the probabilities of searching for different keys in the tree are known, then the tree can be adapted to improve the expected search time. Search trees can also handle insertions and deletions easily, although an unfavorable sequence of such operations may degrade the performance of the tree.

Balanced tree schemes (including AVL trees, 2-3 trees, and B-trees) correct the above-mentioned problem, so that insertions, deletions, and searches can all be performed in logarithmic time in the worst case. Variants of trees have other nice properties, too; they allow merging and splitting of sets, and priority queue operations. B-trees are well-suited to large applications, because they are designed to minimize the number of secondary memory accesses required to perform a search. However, space utilization is not improved by any of these tree schemes, and in fact it is usually increased because of the need for extra pointers.

Hashing is an essentially different approach to the problem. Here a suitable randomizing function is used to compute the location at which a key is stored. Hashing methods are very fast on the average, although the worst case is linear; fortunately this worst case almost never happens.

An interesting variant of hashing, called superimposed coding, was proposed by Bloom [7] (see also [4, §6.5], [8]), and at last provides for reduction in space,
although at the expense of allowing some error. Since this method is perhaps less well known we give a description of it here.

## Superimposed coding

The idea is as follows. We use a single large bit array, initialized to seros, plus a suitable set of $d$ different hash functions. To represent a word, we use the hash functions to compute $d$ bit positions in the large array of bits, and set these bits to ones. We do this for each word in the set. Note that some bits may be set by more than one word.

To test if a word is in the set, we compute the $d$ bit positions asso-iated with the word as above, and check to see if they are all ones in the array. If any of them are zero, the word cannot be in the set, so we reject it. Otherwise if all of the bits are ones, we accept the word. However, some words not in the set might be erroneously accepted, if they happen to hash into bits that are all "covered" by words in the set.

It can be shown [7] that the above scheme makes the best use of space when the density of bits in the array, after all the words have been inserted, is approximately one-half. In this case the probability that a word not in the set is erroneously accepted is $2^{-d}$. For example if each word is hashed into 4 bit positions, the error probability is $1 / 10$. The required size of the bit array is approxinately $n d \lg e$, where $n$ is the number of items in the set, and $\lg e \approx 1.44$.

In fact Bloom specifically discusses automatic hyphenation as an application for his scheme! The scenario is as follows. Suppose we have a relatively compact routine for hyphenation that works correctly for 90 percent of the words in a large dictionary, but it is in error or fails to hyphenate the other 10 percent. We would then like some way to test if a word belongs to the 10 percent, but we do not have room to store all of these words in main memory. If we instead use the superimposed coding scheme to test for these words, the space required can be much reduced. For example with $d=4$ we only need about 6 bits per wcid. The penalty is that some words will be erroncously identified as being in the 10 percent. However, this is acceptable because usually the test word will be rejected and we can then be sure that it is not one of the exceptions. (Either it is in the other 90 percent or it is not in the dictionary at all.) In the comparatively rare case that the word is accepted, we can go to secondary store, to check explicitly if the word is one of the exceptions.

The above technique is actually used in some commercial hyphenation routines. For now, however, $T_{\mathrm{E}} \mathrm{X}$ will not have an external dictionary. Instead we will require that our hyphenation routine be essentially free of error (although it may not achieve complete hyphenation).

An extreme case of superimposed coding should also be mentioned, namely the bit-vector representation of a set. (Imagine that each word is associated with a single bit position, and one bit is allocated for each possible word.) This representation is often very convenient, because it allows set intersection and union to be performed by simple logical operations. But it also requires space proportional to the size of the universe of the set, which is impractical for words longer than three or four characters.

Tries
The final class of data structures we will consider are the digital search trees, first described by de la Briandais [9] and Fredkin [10]. Fredkin also introduced the term "trie" for this class of trees. (The term was derived from the word retrieval, although it is now pronounced "try".)

Tries are distinct from the other data structures discussed so far because they explicitly assume that the keys are a sequence of values over some (finite) alphabet, rather than a single indivisible entity. Thus tries are particularly well-suited for handling variable-length keys. Also, when appropriately implemented, tries can provide compression of the set represented, because common prefixes of words are combined together; words with the same prefix follow the same search path in the trie.

A trie can be thought of as an $m$-ary tree, where $m$ is the number of characters in the alphabet. A search is performed by examining the key one character at a time and using an $m$-way branch to follow the appropriate path in the trie, starting at the root.

We will use the set of 31 most common English words, shown below, to illustrate different ways of implementing a trie.

| A | FOR | IN | THE |
| :--- | :--- | :--- | :--- |
| AND | FROM | IS | THIS |
| ARE | HAD | IT | TO |
| AS | HAVE | NOT | WAS |
| AT | HE | OF | WHICH |
| BE | HER | ON | WITH |
| BUT | HIS | OR | YOU |
| BY | I | THAT |  |

Figure 1. The 91 most common English words.


Figure 5. Linked trie for the 91 most common English words.

Figure 5 shows a linked trie representing this set of words. In a linked trie, the $m$-way branch is performed using a sequential series of comparisons. Thus in Figure 5 each node represents a yes-no test against a particular character. There are two link fields indicating the next node to take depending on the outcome of the test. On a 'yes' answer, we also move to the next character of the key. The underlined characters are terminal nodes, indicated by an extra bit in the node. If the word ends when we are at a terminal node, then the word is in the set.

Note that we do not have to actually store the keys in the trie, because each node implicitly represents a prefix of a word, namely the sequence of characters leading to that node.

A linked trie is somewhat slow because of the sequential testing required for each character of the key. The number of comparisons per character can be as large as $m$, the size of the alphabet. In addition, the two link fields per node are somewhat wasteful of space. (Under certain circumstances, it is possible to eliminate one of these two links. We will explain this later.)

In an indexed trie, the $m$-way branch is performed using an array of size $m$. The elements of the array are pointers indicating the next family of the trie to go to when the given character is scanned, where a "family" corresponds to the group of nodes in a linked trie for testing a particular character of the key. When performing a search in an indexed trie, the appropriate pointer can be accessed by simply indexing from the base of the array. Thus search will be quite fast.

But indexed tries typically waste a lot of space, because most of the arrays have only a few "valid" pointers (for words in the trie), with the rest of the links being null. This is especially common near the bottom of the trie. Figure 6 shows an indexed trie for the set of 31 common words. This representation requires $20 \times 32=$ 832 array locations, compared to 59 nodes for the linked trie.

Various methods have been proposed to remedy the disadvantages of linked and indexed tries. Trabb Pardo [11] describes and analyzes the space requirements of some simple variants of binary tries. Knuth [4, ex. 6.3-20] analyzes a composite method where an indexed tric is used for the first few levels of the trie, switching to sequential search when only a few keys remain in a subtrie. Mehlhorn [12] suggests using a linary search tree to represent each family of a trie. This requires storage proportioual to the number of "valid" links, as in a linked trie, but allows each character of the key to be processed in at most $\log m$ comparisons. Maly [13] has proposed a "compressed trie" that uses an implicit representation to eliminate links entirely. Each level of the trie is represented by a bit array, where the bits indicate whether or not some word in the set passes through the node corresponding to

Figure 6. Indexed trie for the $\$ 1$ most common English words.
that bit. In addition each family contains a field indicating the number of nonsero bits in the array for all nodes to the left of the current family, so that we can find the desired family on the next level. The storage required for each family is thus reduced to $m+\log n$ bits, where $n$ is the total number of keys. However, compressed tries cannot handle insertions and deletions easily, nor do they retain the speed of indexed tries.

## Packed tries

Our idea is to use an indexed trie, but to save the space for null links by packing the different families of the trie into a single large array, so that links from one family may occupy space normally reserved for links for other families that happen to be null. An example of this is illustrated below.

(In the following, we will sometimes refer to families of the indexed trie as states, and pointers as transitions. This is by analogy with the terminology for finite-state machines.)

When performing a search in the trie, we need a way to check if an indexed pointer actually corresponds to the current family, or if it belongs to some other family that just happens to be packed in the same location. This is done by additionally storing the character indexing a transition along with that transition. Thus a transition belongs to a state only if its character matches the character we are indexing on. This test always works if one additional requirement is satisfied, namely that different states may not be packed at the same base location.

The trie can be packed using a first-fit method. That is, we pack the states one at a time, putting each state into the lowest-indexed location in which it will fit (not overlapping any previously packed transitions, nor at an already occupied base location). On numerous examples based on typical word lists, this heuristic works extremely well. In fact, nearly all of the holes in the trie are often filled by transitions from cther states.

Figure 7 shows the result when the indexed trie of Figure 6 is packed into a single array using the first-fit method. (Actually we have used an additional compression technique called suffix compression before packing the trie; this will be explained in the next section.) The resulting tric fits into just 60 locations. Note


Figure 7. Packed trie for the 31 most common English words.
that the packed trie is a single large array; the rows in the figure should be viewed as one long row.

As an example, here's what happens when we search for the word HAVE in the parked trie. We associate the values 1 through 26 with the letters A through $Z$. The root of the trie is packed at location 0 , so we begin by looking at location 8 corresponding to the letter H. Since ' $\mathrm{H} 3 \mathrm{O}^{\prime}$ ' is stored there, this is a valid transition and we then go to location 30 . Indexing by the letter $A$, we look in location 31, which tells us to go to 29. Now indexing by V gets location 51 , which points to 2. Finally indexing by E gets location 7, which is underlined, indicating that the word HAVE is indeed in the set.

## Suffix compression

A big advantage of the trie data structure is that common prefixes of words are combined automatically into common paths in the trie. This provides a good deal of compression. To save more space, we can try to take advantage of common suffixes.

One way of doing this is to construct a trie in the usual manner, and then merge common subtries together, starting from the leaves (lieves) and working upward. We call this process suffix compression.

For example, in the linked trie of Figure 5 the terminal nodes for the words HIS and THIS, both of which test for the letter S and have no successors, can be combined into a single node. That is, we can let their parent nodes both point to the same node; this does not change the set of words accepted by the trie. It turns out that we can then combine the parent nodes, since both of them test for I and-go to the S node if successful, otherwise stop (no left successor). However, the grandparent nodes (which are actually siblings of the I nodes) cannot be combined even though they both test for $E$, because one of them goes to a terminal $R$ node upon success, while the other has no right successor.

With a larger set of words, a great deal of merging can be possible. Clearly all leaf nodes (nodes with no successors) that test the same character can be combined together. This alone saves a number of nodes equal to the number of words in the dictionary, minus the number of words that are prefixes of other words, plus at most 26. In addition, as we might expect, longer suffixes such as -1 y , -ing , or -tion can frequently be combined.

The suffix compression process may sound complicated, but actually it can be described by a simple recursive algorithm. For each node of the trie, we first compress each of its subtries, then determine if the node can be merged with some other node. In effect, we traverse the trie in depth-first order, checking each node to see if it is equivalent to any previously seen node. A hash table can be used to identify equivalent nodes, based on their (merged) transitions.

The identification of nodes is somewhat easicr using a binary tree representation of the trie, rather than an $m$-ary representation, because each node will then have just two link fields in addition to the character and output bit. Thus it will be convenient to use a tinked trie when performing suffix compression. The linked representation is also more convenient for constructing the trie in the first place, because of the ease of performing insertions.

After applying suffix compression, the trie can be converted to an indexed trie and packed as described previously. (We should remark that performing suffix compression on a linked trie can yield some additiona' :ompression, because trie families can be partially merged. However such compression is lost when the trie is converted to indexed form.)

The author has performed numerous experiments with the above ideas. The results for some representative word lists are shown in Table 1 below. The last three
columns show the number of nodes in the linked, suffix-compressed, and packed tries, respectively. Each transition of the packed trie consists of a pointer, a character, and a bit indicating if this is an accepting transition.

| word list | words | characters | linked | compressed | packed |
| :--- | ---: | ---: | ---: | ---: | ---: |
| pascal | 35 | 145 | 125 | 104 | 120 |
| murray | 2720 | 19,144 | 8039 | 4272 | 4285 |
| pocket | 31,036 | 247,612 | 92,339 | 38,619 | 38,638 |
| unabrd | 235,545 | $2,250,805$ | 759,045 | - | - |

## Table 1. Sufix-compressed pucked tries.

The algorithms for building a linked trie, suffix compression, and first-fit packing are used in $\mathrm{T}_{\mathrm{E}} \mathrm{X} 82$ to preprocess the set of hyphenation patterns into a packed trie used by the hyphenation routine. A WEB description of these algorithms can be found in [14].

## Derived forms

Most dictionaries do not list the most common derived forms of words, namely regular plurals of nouns and verbs ( -s forms), participles and gerunds of verbs (-ed and -ing forms), and comparatives and superlatives of adjectives (-er and -est). This makes sense, because a user of the dictionary can easily determine when a word possesses one of these regular forms. However, if we use the word list from a typical dictionary for spelling checking, we will be faced with the problem of determining when a word is one of these derived forms.

Some spelling checkers deal with this problem by attempting to recognize affixes. This is done not only for the derived forms mentioned above but other common variant forms as well, with the purpose of reducing the number of words that have to be stored in the dictionary. A set of logical rules is used to determine when certain prefixes and suffixes can be stripped from the word under consideration.

However such rules can be quite complicated, and they inevitably make errors. The situation is not unlike that of finding rules for hyphenation, which should not be surprising, since affix recognition is an important part of any rule-based hyphenation algorithm. This problem has been studied in some detail in a series of papers by Resnikoff and Dolby [15].

Since affix recognition is difficult, it is preferable to base a spelling checker on a complete word list, including all derived forms. However, a lot of additional space will be required to store all of these forms, even though much of the added data is
redundant. We might hope that some appropriate method could provide substantial compression of the expanded word list. It turns out that suffix-compressed tries handle this quite well. When derived forms were added to our pocket dictionary word list, it increased in size to 40,858 words and 404,916 characters, but the resulting packed trie only increased to 46,553 transitions (compare the pocket dictionary statistics in Table 1).
-Hyphenation programs also need to deal with the problem of derived forms. In our pattern-matching approach, we intend to extract the hyphenation rules automatically from the dictionary. Thus it is again preferable for our word list to include all derived forms.

The creation of such an expanded word list required a good deal of work. The author had access to a computer-readable copy of Webster's Pocket Dictionary [16], including parts of speech and definitions. This made it feasible to identify nouns, verbs, etc., and to generate the appropriate derived forms mechanically. Unfortunately the resulting word lists required extensive editing to eliminate many never-used or somewhat nonsensical derived forms, e.g. 'informations'.

## Spelling checkers

Computer-based word processing systens iave recently come into widespread use. As a result there has been a surge of interest in programs for automatic spelling checking and correction. Here we will consider the dictionary representations used by some existing spelling checkers.

One of the earliest programs, designed for a large timesharing computer, was the DEC-10 SPELL program written by Ralph Gorin [17]. It uses a 12,000 word dictionary stored in main memory. A simple hash function assigns a unique 'bucket' to each word depending on its length and the first two characters. Words in the same bucket are listed sequentially. The number of words in each bucket is relatively small (typically 5 to 50 words), so this representation is fairly efficient for searching. In addition, the buckets provide convenient access to groups of similar words; this is useful when the program tries to correct spelling errors.

The dictionary used by SPELL does not contain derived forms. Instead some simple affix stripping rules are normally used; the author of the program notes that these are "error-prone".

Another spelling checker is described by James i. Peterson [18]. His program uses three separate dictionarics: (1) a small list of 258 common English words, (2) a dynamic 'cache' of about 1000 document-specific words, and (3) a large, comprehensive dictionary, stored on disk. The list of common words (which is static) is represented using a suffix-compressed linked trie. The dynamic cache is maintained
using a hash table. Both of these dictionaries are kept in main memory for speed. The disk dictionary uses an in-core index, so that at most one disk access is required per search.

Robert Nix [19] describes a spelling checker based on the superimposed coding method. He reports that this method allows the dictionary from the SPELL program to be compressed to just 20 percent of its original size, while allowing $0.1 \%$ cbance of error.
$\AA$ considerably different approach to spelling checking was taken by the TYPO program developed at Bell Labs [20]. This program uses digram and trigram frequencies to identify "improbable" words. After processing a document, the words are listed in order of decreasing improbability for the user to peruse. (Words appearing in a list of 2726 common technical words are not shown.) The authors report that this format is "psychologically rewarding", because many errors are found at the beginning, inducing the user to continue scanning the list until errors become rare.

In addition to the above, there have receritly been a number of spelling checkers developed for the "personal computer" market. Because these programs run on small microprocessor-based systems, it is especially important to reduce the size of the dictionary. Standard techniques include hash coding (allowing some error), incore caches of common words, and special codes for common prefixes and suffixes. One program first constructs a sorted list of all words in the document, and then compares this list with the dictionary in a single sequential pass. The dictionary can then be stored in a compact form suited for sequential scanning, where each word is represented by its difference from the previous word.

Besides simply detecting when words are not in a dictionary, the design of a practical spelling checker involves a number of other issues. For example many spelling checkers also try to perform spelling correction. This is usually done by searching the dictionary for words similar to the misspelled word. Errors and suggested replacements can be presented in an interactive fashion, allowing the user to see the context from the document and make the necessary changes. The contents of the dictionary are of course very important, and each user may want to modify the word list to match his or her own vocabulary. Finally, a plain spelling checker cannot detect problems such as incorrect word usage or mistakes in grammar; a more sophisticated program performing syntactic and perhaps semantic analysis of the text would be necessary.

Conclusion and related ideas
The dictionary problem is a fundamental problem of computer science, and it has many applications besides spelling checking. Most data structures for this problem consider the elements of the set as atomic entities, fitting into a single computer word. However in many applications, particularly word processing, the keys are actually variable-length strings of characters. Most of the standard techniques are somewhat awkward when dealing with variable length keys. Only the trie data structure is well-suited for this situation.

We have proposed a variant of tries that we call a packed trie. Search in a packed trie is performed by indexing, and it is therefore very fast. The first-fit packing technique usually produces a fairly compact representation as well.

We have not djscussed how to perform dynamic insertions and deletions with a packed trie. In Chapter 4 we discuss a way to handle this problem, when no suffix compression is used, by repacking states when necessary.

The idea of suffix compression is not new. As mentioned, Peterson's spelling checker uses this idea also. But in fact, if we view our trie as a finite-state machine, suffix compression is equivalent to the well-known idea of state minimization. In our case the machine is acyclic, that is, it has no loops.

Suffix compression is also closely related to the common subexpression problem from compiler theory. In particular, it can be considered a special case of a problem called acyclic congruence closure, which has bcen studied by Downey, Sethi, and Tarjan [21]. They give a linear-time algorithm for suffix compression that does not use hashing, but it is somewhat complicated to implement and requires additional data structures.

The idea for the first-fit packing method was inspired by the paper "Storing a sparse table" by Tarjan and Yao [22]. The technique has been used for compressing parsing tables, as discussed by Zeigler [23] (see also [24]). However, our packed trie implementation differs somewhat from the applications discussed in the above references, because of our emphasis on space minimization. In particular, the idea of storing the character that indexes a transition, along with that transition, seems to be new. This has an advantage over other techniques for distinguishing states, such as the use of back pointers, because the character requires fewer bits.

The paper by Tarjan and Yao also contains an interesting theorem characterizing the performance of the first-fit packing method. They consider a modification suggested by Zeigler, where the states are first sorted into decreasing order based on the number of non-null transitions in each state. The idea is that small states, which can be packed more easily, will be saved to the end. They prove that if the
distribution of transitions among states satisfies a "harmonic decay" condition, then essentially all of the holes in the first-fit packing will be filled.

More precisely, let $n(l)$ be the total number of non-null transitions in states with more than $l$ transitions, for $l \geq 0$. If the harmonic decay property $n(l) \leq n /(l+1)$ is satisfied, then the first-fit-decreasing packing satisfies $0 \leq b(i) \leq n$ for all $i$, where $n=n(0)$ is the total number of transitions and $b(i)$ is the base location at which the $i$ th state is packed.

The above theorem does not take into account our additional restriction that no two states may be packed at the same base location. When the proof is modified to include this restriction, the bound goes up by a factor of two. However in practice we seem to be able to do much better.

The main reason for the good performance of the first-fit packing scheme is the fact that there are usually enough single-transition states to fill in the holes created by larger states. It is not really necessary to sort the states by number of transitions; any packing order that distributes large and small states fairly evenly will work well. We have found it convenient simply to use the order obtained by traversing the linked trie.

Improvements on the algorithms discussed in this chapter are possible in certain cases. If we store a linked trie in aspecific traversal order, we can eliminate one of the link fields. For example, if we list the nodes of the trie in preorder, the left successor of a node will always appear immediately after that node. An extra bit is used to indicate that a node has no left successor. Of course this technique works for other types of trees as well.

If the word list is already sorted, linked trie insertion can be performed with only a small portion of the tric in memory at any time, namely the portion along the current insertion path. This can be a great advantage if we are are processing a large dictionary and cannot store the entire linked trie in memory.

# Hyphenation 

Let us now try to apply the ideas of the previous chapter to the problem of hyphenation. $\mathrm{T}_{\mathrm{E}} \mathrm{X} 82$ will use the pattern matching method described in Chapter 1, but we shall first discuss some related approaches that were considered.

## Finite-state machines with output

We can modify our trie-based dictionary representation to perform hyphenation by changing the output of the trie (or finite-state machine) to a multiple-valued output indicating how the word can be hyphenated, instead of just a binary yes-no output indicating whether or not the word is in the dictionary. That is, instead of associating a single bit with each trie transition, we would have a larger "output" field indicating the hyphenation "action" to be taken on this transition. Thus on recognizing the word hy-phen-a-tion, the output would say "you can hyphenate this word after the second, sixth, or seventh letters".

To represent the hyphenation output, we could simply list the hyphen positions, or we could use a bit vector indicating the allowable hyphen points. Since there are only a few hundred different outputs and most of them occur many times, we can save some space by assigning each output a unique code and storing the actual hyphen positions in a separate table.

To conveniently handle the variable number of hyphen positions in outputs, we will use a linked representation that allows different outputs to share common portions of their output lists. This is implemented using a hash table containing pairs of the form (output, next), where output is a hyphenation position and next is a (possibly null) pointer to another entry in the table. To add a new output list to the table, we hash each of its outputs in turn, making each output point to the previous one. Interestingly, this process is quite similar to suffix compression.

The trie with hyphenation output can be suffix-compressed and packed in the same manner as discussed in Chapter 2. Because of the greater variety of outputs more of the subtries will be distinct, and there is somewhat less compression.

From our pocket dictionary (with hyphens), for example, we obtained a packed trie occupying 51,699 locations.

We can improve things slightly by "pushing outputs forward". That is, we can output partial hyphenations as soon as possible instead of waiting until the end of the word. This allows some additional suffix compression.

For example, upon scanning the letters hyph at the beginning of a word, we can already say "hyphenate after the second letter" because this is allowed for all words beginning with those letters. Note we could not say this after scanning j. st hyp, because of words like hyp-not-ic. Upon further scanning ena, we can say "hyphenate after the sixth letter".

When implementing this idea, we run into a small problem. There are quite a few words that are prefixes of other words, but hyphenate differently on the letters they have in common, e.g. ca-ret and care-tak-er, or as-pi-rin and as-pir-ing. To avoid losing hyphenation output, we could have a separate output whenever an end-of-word bit appears, but a simpler method is to append an end-ofword character to each word before inserting it into the trie. This increases the size of the linked trie considerably, but suffix compression merges most of these nodes together.

With the above modifications, the packed trie for the pocket dictionary was reduced to 44,128 transitions.

Although we have obtained substantial compression of the dictionary, the result is still too large for our purposes. The problem is that as long as we insist that only words in the dictionary be hyphenated, we cannot hope to reduce the space required to below that needed for spelling checking alone. So we must give up this restriction.

For example, we could eliminate the end-of-word bit. Then after pushing outputs forward, we can prune branches of the trie for which there is no further output. This would reduce the pocket dictionary trie to 35,429 transitions.

## Minimization with don't cares

In this sectioli we describe a more drastic approach to compression that takes advantage of situations where we "don't care" what the algorithm dops.

As previously noted, most of the states in an indexed trie are quite sparse; that is, only a few of the characters have explicit transitions. Since the missing transitions are never accessed by words in our dictionary, we can allow them to be filled by arbitrary transitions.

This should not be confused with the overlapping of states that may occur in the trie-packing process. Instead, we mean that the added transitions will actually become part of the state.

There are two ways in which this might allow us to save more space in the minimization process. First, states no longer have to be identical in order to be merged; they only have to agree on those characters where both (or all) have explicit transitions. Second, the merging of non-equivalent states may allow further merging that was not previously possible, because some transitions have now become equivalent.

For example, consider again the trie of Figure 5. When discussing suffix compression, we noted that the terminal S nodes for the words HIS and THIS could be merged together, but that the parent chains, each containing transitions for A, E, and I, could not be completely merged. However, in minimization with don't cares these two states can be merged. Note that such a merge will require that the DV state below the first A be merged with the $T$ below the second $A$; this can be done because those states have no overlapping transitions.

As another example, notice that if the word AN were added to our vocabulary, then the NRST chain succeeding the root A node could be merged with the NST chain below the initial I node. (Actually, it doesn't make much sense to do minimization with don't cares on a trie used to recognize words in a dictionary, but we will ignore that objection for the purposes of this example.)

Unfortunately, trie minimization with don't cares seems more complicated than the suffix-compression process of Chapter 2. The problem is that states can be merged in more than one way. That is, the collection of mergeable states no longer forms an equivalence relation, as in regular finite-state minimization. In fact, we can sometimes obtain additional compression by allowing the same state to appear more than once. Another complication is that don't care merges can introduce loops into our trie.

Thus it seems that finding the minimum size trie will be difficult. Pfleeger [25] has shown this problem to be NP-complete, by transformation from graph coloring; however, his construction requires the number of transitions per state to be unbounded. It may be possible to remove this requirement, but we have not proved this.

So in order to experiment with trie minimization with don't cares, we have made some simplifications. We start by performing suffix compression in the usual manner. We then go through the states in a bottom-up order, checking each to see if it can be merged with any previous state by taking advantage of don't cares. Note that such merges may require further merges among states already seen.

We only try merges that actually save space, that is, where explicit transitions are merged. Otherwise, states with only a few transitions are very likely to be mergeable, but such merges may constrain us unnecessarily at a later stage of the minimization. In addition, we will not consider having multiple copies of states.

Even this simplified algorithm can be quite time consuming, so we did not try it on our pocket dictionary. On a list of 2726 technical words, don't care minimization reduced the number of states in the suffix-compressed, output-pruned trie from 1685 to just 283, while the number of transitions was reduced from 3627 to 2427. However, because the resulting states were larger, the first-fit packing performed rather poorly, producing a packed trie with 3408 transitions. So in this case don't care minimization yielded an additional compression of less than 10 percent.

Also, the behavior of the resulting hyphenation algorithm on words not in the dictionary became rather unpredictable. Once a word leaves the "known" paths of the packed trie, strange things might happen!

We can get even wilder effects by carrying the don't care assumption one step further, and eliminating the character field from the packed trie altogether (leaving just the output and trie link). Words in the dictionary will always index the correct transitions, but on other words we now have no way of telling when we have reached an invalid trie transition.

It turns out that the problem of state minimization with don't cares was studied in the 1960s by electrical engineers, who called it "minimization of incompletely specified sequential machines" (see e.g. [26]). However, typical instances of the problem involved machines with only a few states, rather than thousands as in our case, so it was often possible to find a minimized machine by hand. Also, the emphasis was on minimizing the number of states of the machine, rather than the number of state transitions.

In ordinary finite-state minimization, these are equivalent, but don't care minimization can actually introduce extra transitions, for example when states are duplicated. In the old days, finite-state machines were implemented using combinational logic, so the most important consideration was to reduce the number of states. In our trie representation, however, the space used is proportional to the number of transitions. Furthermore, finite-state machines are now often implemented using PLA's (programmed logic arrays), for which the number of transitions is also the best measure of space.

## Pattern matching

Since trie minimization with don't cares still doesn't provide sufficient compression, and since it lead to unpredictable behavior on words not in the dictionary,
we need a different approach. It seems rxpensive to insist on complete hyphenation of the dictionary, so we will give up this requirement. We could allow some errors; or to be safer, we could allow some hyphens to be missed.

We now return to the pattern matching apprcach described in Chapter 1. Some further arguments as to why this method seems advantageous are given below. We should first reassure the reader that all the discussion so far has not been in vain, because a packed trie will be an ideal data structure for representing the patterns in the final hyphenation algorithm. Here the outputs will include the hyphenation level as well as the intercharacter position.

Hyphenating and inhibiting patterns allow considerable flexibility in the performance of the resulting algorithm. For example, we could allow a certain amount of error by using patterns that aren't always safe (but that presumably do find many correct hyphens).

We can also restrict ourselves to partial hyphenation in a natural way. That is, it turns out that a relatively small number of patterns will get a large fraction of the hyphens in the dictionary. The remaining hyphens become harder and harder to find, as we are left with mostly exceptional cases. Thus we can choose the most effective patierns first, taking more and more specialized patterns until we run out of space.

In addition, patterns perform quite well on words not in the dictionary, if those words follow "normal" pronunciation rules.

Patterns are "context-free"; that is, they can apply anywhere in a word. This seems to be an important advantage. In the trie-based approach discussed earlier in this chapter, a word is always scanned from beginning to end and each state of the trie 'remembers' the entire prefix of the word scanned so far, even if the letters scanned near the beginning no longer affect the hyphenation of the word. Suffix compression eliminates some of this unnecessary state information, by combining states that are identical with respect to future hyphenation. Minimization with don't cares takes this further, allowing 'similar' states to be combined as long as they behave identically on all characters that they have in common.

However, we have seen that it is difficult to guide the minimization with don't cares to achieve these reductions. Patterns embody such don't care situations naturally (if we can find a good way of selecting the patterns).

The context-free nature of patterns helps in another way, as explained below. Recall that we will use a packed trie to represent the patterns. To find all patterns that match in a given word, we perform a search starting at each letter of the word. Thus after completing a search starting from some letter position, we may have to
back up in the word to start the next search. By contrast, our original trio-based approach works with no backup.

Suppose we wanted to convert the pattern trie into a finite-state recognizer that works with no backup. This can be done in two stages. We first add "failure links ${ }^{n}$ to each state that tell which state to go to if there is no explicit transition for the current character of the word. The failure state is the state in the trie that we would have reached, if we had started the search one letter later in the word.

Next, we can convert the failure-link machine into a true finite-state machine by.filling in the missing transitions of each state with those of its failure state. (For more details of this process, see [27], [28].)

However, the above state merging will introduce a lot of additional transitions. Even using failure links requires one additional pointer per state. Thus by performing pattern matching with backup, we seem to save a good deal of space. And in practice, long backups rarely occur.

Finally, the idea of inhibiting patterns seems to be very useful. Such patterns extend the power of a finite-state machine, somewhat like adding the "not" operator to regular expressions.

## Pattern generation

We now discuss how to choose a suitable set of patterns for hyphenation. In order to decide which patterns are "good", we must first specify the desired properties of the resulting hyphenation aigorithm.

We obviously want to maximize the number of hyphens found, minimize the error, and minimize the space required by our algorithm. For example, we could try to maximize some (say linear) function of the above three quantities, or we could hold one or two of the quantities constant and optimize the others.

For TEX82, we wanted a hyphenation algorithmimeeting the following requirements. The algorithm should use only a moderate amount of space ( $20-30 \mathrm{~K}$ bytes), including any exception dictionary; and it should find as many hyphens as possible, while making little or no error. This is similar to the specifications for the original $T_{E} X$ algorithm, except that we now hope to find substantially more hyphens.

Of course, the results will depend on the word list used. We decided to base the algorithm on our copy of Webster's Pocket Dictionary, mainly because this was the only word list we had that included all derived forms.

We also thought that a larger dictionary would contain many rare or specialized words that we might not want to worry about. In $\mathrm{P}^{-}$ticular, we did not want such infrequent words to affect the choice of patterns, because we hoped to obtain a set of patterns embodying many of the "usual" rules for hyphenation.

In developing the TEX82 algorithm, however, the word list was tuned up considerably. $\Lambda$ few thousand common words were weighted more heavily so that they would be more likely to be hyphenated. In fact, the current algorithm guarantees complete hyphenation of the 676 most common English words (according to [29]), as well as a short list of common technical words (e.g. al-go-rithm).

In addition, over 1000 "exception" words have been added to the dictionary, to ensure that they would not be incorrectly hyphenated. Most of these were found by testing the algorithm (based on the initial word list) against a larger dictionary obtained from a publisher, containing about 115,000 entries. This produced about

10,000 errors on words not in the pocket dictionary. Most of these were specialized technical terms that we decided not to worry about, but a few hundred were embarrassing enough that we decided to add them to the word list. These included compound words (camp-fire), proper names (Af-ghan-i-stan), and new words (bio-rhythm) that probably did not exist in 1966, when our pocket dictionary was originally put online.

After the word list was augmented, a new set of patterns was generated, and a new list of exceptions was found and added to the list. Fortunately this process seemed to converge after a few iterations.

## Heuristics

The selection of patterns in an 'optimal' way seems very difficult. The problem is that several patterns may apply to a particular hyphen point, including both hyphenating and inhibiting patterns. Thus complicated interactions can arise if we try to determine, say, the minimum set of patterns finding a given number of hyphens. (The situation is somewhat analogous to a set cover problem.)

Instead, we will select patterns in a series of "passes" through the word list. In each pass we take into account only the effects of patterns chosen in previous passes. Thus we sidestep the problem of interactions mentioned above.

In addition, we will define a measure of pattern "efficiency" so that we can use a greedy approach in each pass, selecting the most efficient patterns.

Patterns will be selected one level at a time, starting with a level of hyphenating patterns. Patterns at each level will be selected in order of increasing pattern length.

Furthermore patterns of a given length applying to different intercharacter positions (for example -tio and t-io) will be selected in separate passes through the dictionary. Thus the patterns of length $n$ at a given level will be chosen in $n+1$ passes through the dictionary.

At first we did not do this, but selected all patterns of a given length (at a given level) in a single pass, to save time. However, we found that this resulted in considerable duplication of effort, as many hyphens were covered by two or more patterns. By considering different intercharacter positions in separate passes, there is never any overlap among the patterns selected in a single pass.

In each pass, we collect statistics on all patterns appearing in the dictionary, counting the number of times we could hyphenate at a particular point in the pattern, and the number of times we could not.

For example, the pattern tio appears 1793 times in the pocket dictionary, and in 1773 cases we can hyphenate the word before the $t$, while in 20 cases we can
not. (We only count instances where the hyphen position occurs at least two letters from either edge of the word.)

These counts are used to determine the efficiency rating of patterns. For example if we are considering only "safe" patterns, that is, paiterns that can always be hyphenated at a particular position, then a reasonable rating is simply the number of hyphens found. We could then decide to take, say, all patterns finding at least a given number of hyphens.

However, most of the patterns we use will make some error. How should these patterns be evaluated? In the worst case, errors can be handled by simply listing them in an exception dictionary. Assuming that one unit of space is required to represent each pattern as well as each exception, the "efficiency" of a pattern could be defined as eff $=\operatorname{good} /(1+b a d)$ where good is the number of hyphens correctly found and bad is the number of errors made.
(The space used by the final algorithm really depends on how much compression is produced by the packed trie used to represent the patterns, but since it is hard to predict the exact number of transitions required, we just use the number of patterns as an approximate measure of size.)

By using inhibiting patterns, however, we can often do better than listing the exceptions individually. The quantity bad in the above formula should then be devalued a bit depending on how effective patterns at the next level are. So a better formula might be

$$
e f f=\frac{\text { good }}{1+\text { bad } / b a d_{-} e f f},
$$

where bad_eff is the estimated efficiency of patterns at the next level (inhibiting errors at the current level).

Note that it may be difficult to determine the efficiency at the next level, when we are still deciding what patterns to take at the current level! We will use a pattern selection criterion of the form eff $\geq$ thresh, but we cannot predict exactly how many patterns will be chosen and what their overall performance will be. The best we can do is use reasonable estimates 'ased on previous runs of the pattern generation program. Some statistics from trial runs of this program are presented later in this chapter.

## Collecting pattern statistics

So the main task of the pattern generation process is to collect count statistics about patterns in the dictionary. Because of time and space limitations this becomes an interesting data structure exercise.

For short (length 2 and 3) patterns, we can simply use a table of size $26^{2}$ or $26^{3}$, respectively, to hold the counts during a pass through the dictionary. For longer patterns, this is impractical.

Here's the first approach we used for longer patterns. In a pass through the dictionary, every occurrence of a pattern is written out to a file, along with an indication of whether or not a hyphen was allowed at the position under consideration, The file of patterns is sorted to bring identical patterns together, and then a pass is made through the sorted list to compile the count statistics for each pattern.

This approach makes it feasible to collect statistics for longer length patterns, and was used to conduct our initial experiments with pattern generation. However it is still quite time and space consuming, especially when sorting the large lists of patterns. Note that an external sorting algorithm is usually necessary.

Since only a fraction of the possible patterns of a particular length actually occur in the dictinnary, we could instead store them in a hash table or one of the other data structures discussed in Chapter 2. It turns out that a modification of our packed trie data structure is well-suited to this task. The advantages of the packed trie are very fast lookup, compactness, and graceful handling of variable length patterns.

Combined with some judicious "pruning" of the patterns that are considered, the memory requirements are much reduced, allowing the entire pattern selection process to be carried out "in core" on our PDP-10 computer,

By "pruning" patterns we mean the following. If a pattern contains a shorter pattern at the same level that has already been chosen, the longer pattern obviously need not be considered, so we do not have to count its occurrences. Similarly, if a pattern appears so few times in the dictionary thet under the current selection criterion it can never be chosen, then we can mark the pattern as "hopeless" so that any longer patterns at this level containing it need not be considered.

Pruning greatly reduces the number of patterns that must be considered, especially at longer lengths.
Dynamic packed tries
Unlike the static dictionary problem considered in Chapter 2, the set of patterns to be represented is not known in advance. In order to use a packed trie for storing the patterns being considered in a pass through the dictionary, we need some way to dynamically insert new patterns into the trie.

For any pattern, we start by performing a search in the packed trie as usual, following existing links until reaching a state where a new trie transition must be
added. If we are lucky, the location needed by the new transition will still be empty in the packed trie, otherwise we will have to do some repacking.

Note that we will not be using suffix compression, because this complicates things considerably. We would need back pointers or reference counts to determine what nodes need to be unmerged, and we would need a hash table or other auxiliary information in order to remerge the newly added nodes. Furthermore, suffix merging does not produce a great deal of compression on the relatively short patterns we will be dealing with.

The simplest way of resolving the packing conflict caused by the addition of a new transition is to just repack the changed state (and update the link of its parent state). To maintain good space utilization, we should try to fit the modified state among the holes in the trie. This can be done by maintaining a dynamic list of unoccupied cells in the trie, and using a first-fit search.

However, repacking turns out to be rather expensive for large states that are unlikely to fit into the holes in the trie, unless the array is very sparse. We can avoid this by packing such states into the free space immediately to the right of the occupied locations. The size threshold for attempting a first-fit packing can be adjusted depending on the density of the array, how much time we are willing to spend on insertions, or how close we are to running out of room.

After adding the critical transition as discussed above, we may need to add some more trie nodes for the remaining characters of the new pattern. These new states contain just a single transition, so they should be easy to fit into the trie.

The pattern generation program uses a second packed trie to store the set of patterns selected so far. Recall that, before collecting statistics about the patterns in each word, we must first hyphenate the word according to the patterns chosen in previous passes. This is done not only to determine the current partial hyphenation, but also to identify pruned patterns that need not be considered. Once again, the advantages of the packed trie are compactness and very fast "hyphenation".

At the end of a pass, we need to add new patterns, including "hopeless" patterns, to the trie. Thus it will be convenient to use a dynamic packed trie here as well. At the end of a level, we probably want to delete hopeless patterns from the trie in order to recover their space, if we are going to generate more levels. This turns out to be relatively easy; we just remove the appropriate output and return any freed nodes to the available list.

Below we give some statistics that will give an idea of how well a dynamic packed trie performs. We took the current set of 4447 hyphenation patterns, randomized them, and then inserted them one-by-one into a dynamic packed trie.
(Note that in the situations described above, there will actually be many searches per insertion, so we can afford some extra effort when performing insertions.) The patterns occupy 7214 trie nodes, but the packed trie will use more locations, depending on the setting of the first-fit packing threshold. The columns of the table show, respectively, the maximum state size for which a first-fit packing is attempted, the number of states packed, the number of locations tried by the first-fit procedure (this dominates the running time), the number of states repacked, and the number of locations used in the final packed trie.

| thresh | pack | first_fit | unpack | trie_max |
| :---: | :---: | :---: | :---: | :---: |
| $\boldsymbol{\infty}$ | 6113 | 877,301 | 2781 | 9671 |
| 13 | 6060 | 761,228 | 2728 | 9458 |
| 9 | 6074 | 559,835 | 2742 | 9606 |
| $\mathbf{7}$ | 6027 | 359,537 | 2695 | 9606 |
| 5 | 5863 | 147,468 | 2531 | 10,366 |
| 4 | 5746 | 63,181 | 2414 | 11,209 |
| 3 | 5563 | 33,826 | 2231 | 13,296 |
| 2 | 5242 | 10,885 | 1010 | 15,009 |
| 1 | 4847 | 8956 | 1515 | 16,536 |
| $\mathbf{0}$ | 4577 | 6073 | 1245 | $\mathbf{1 8 , 6 2 8}$ |

Table 2. Dynamic packed trie statistice.

## Experimental results

We now give some results from trial runs of the pattern generation program, ${ }^{*}$ and explain how the current TEX82 patterns were generated. As mentioned earlier, the development of these patterns involved some augmentation of the word list. The results described here are based on the latest version of the dietionary.

At each level, the selection of patterns is controlled by three parameters called good_wt, bad_wt, and thresh. If a pattern can be hyphenated good times at a particular position, but makes bad errors, then it will be selected if

$$
\text { good } * \text { good_wt - bad } * \text { bad_wt } \geq \text { thresh. }
$$

Note that the efficiency formula given earlier in this chapter can be converted into the above form.

We can first try using only safe patterns, that is, patterns that can always be hyphenated at a particular position. The table below shows the results when all safe patterns finding at least a given number of hyphens are chosen. Note that

| parameters | patterns | hyphens | percent |
| :---: | ---: | :---: | :---: |
| $1 \propto 40$ | 401 | 31,083 | $35.2 \%$ |
| $1 \infty 20$ | 1024 | $45,310 \ldots$ | $51.3 \%$ |
| $1 \infty 10$ | 2272 | 58,580 | $66.3 \%$ |
| $1 \infty 5$ | 4603 | 70,014 | $79.2 \%$ |
| $1 \infty 3$ | 7052 | 76,236 | $86.2 \%$ |
| $1 \infty 2$ | 10,456 | 83,450 | $94.4 \%$ |
| $1 \infty 1$ | 16,336 | 87,271 | $98.7 \%$ |

Table 3. Safe hyphenating patterns.
an infinite bad_wt ensures that only safe patterns are chosen. The table shows the number of patterns obtained, and the number and percentage of hyphens found.

We see that, roughly speaking, halving the threshold doubles the number of patterns, but only increases the percentage of hyphens by a constant amount. The last 20 percent or so of hyphens become quite expensive to find.
(In order to save computer time, we have only considered patterns of length 6 or less in obtaining the above statistics, so the figures do not quite represent all patterns above a given threshold. In particular, the patterns at threshold 1 do not find $100 \%$ of the hyphens, although even with indefinitely long patterns there would still be a few hyphens that would not be found, such as re-cord.)

The space required to represent patterns in the final algorithm is slightly more than one trie transition per pattern. Each transition occupies 4 bytes ( 1 byte each for character and output, plus 2 bytes for trie link). The output table requires an additional 3 bytes per entry (hyphenation position, value, and next output), but there are only a few hundred outputs. Thus to stay within the desired space limitations for $\mathrm{T}_{\mathrm{E}} \mathrm{X} 82$, we can use at most about 5000 patterns.

We next try using two levels of patterns, to see if the idea of inhibiting patterns actually pays off. The results are shown below, where in each case the initial level of hyphenating patterns is followed by a level of inhibiting patterns that remove nearly all of the error.

The last set of patterns achicves $86.7 \%$ hyphenation using 4696 patterns. By contrast, the $1 \infty 3$ patterns from the previous table achieves $86.2 \%$ with 7052 patterns. So inhibiting patterns do help. In addition, notice that we have only used "safe" inhibiting patterns above; this means that none of the good hyphens are lost. We can do better by using patterns that also inhibit some correct hyphens.

After a good deal of further experimentation, we decided to use five levels of patterns in the current TEX82 algorithm. The reason for this is as follows. In

|  | parameters | patterns | hyphens |  | percent |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 12020 | 816 | 51,359 | 505 | 58.1\% | 0.6\% |
|  | $1 \infty 1$ | 315 | 0 | 463 | 58.1\% | 0.1\% |
|  | 11010 | 1510 | 64,893 | 1694 | 73.5\% | 1.9\% |
|  | $1 \infty 1$ | 824 | 0 | 1531 | 73.5\% | 0.2\% |
|  | 155 | 2573 | 76,632 | 5254 | 86.7\% | 5.9\% |
| $\cdots$ | $1 \infty 1$ | 2123 | 0 | 4826 | 86.7\% | 0.5\% |

Table 4. Two levels of patterns.
addition to finding a high percentage of hyphens, we also wanted a certain amount of guaranteed behavior. That is, we wanted to make essentially no errors on words in the dictionary, and also to ensure complete hyphenation of certain common words.

To accomplish this, we use a final level of safe hyphenating patterns, with the threshold set as low as feasible (in our case 4). If we then weight the list of important words by a factor of at least 4 , the patterns obtained will hyphenate them completely (except when a word can be hyphenated in two different ways).

To guarantee no error, the level of inhibiting patterns immediately preceding the final level should have a threshold of 1 so that even patterns applying to a single word will be chosen. Note these do not need to be "safe" inhibiting patterns, since the final level will pick up all hyphens that should be found.

The problem is, if there are too many errors remaining before the last inhibiting level, we will need too many patterns to handle them. If we use three levels in all, then the initial level of hyphenating patterns can allow just a small amount of error.

However, we would like to take advantage of the high efficiency of hyphenating patterns that allow a greater percentage of error. So instead, we will use an initial level of hyphenating patterns with relatively high threshold and allowing considerable error, followed by a 'coarse' level of inhibiting patterns removing most of the initial error. The third level will consist of relatively safe hyphenating patterns with a somewhat lower threshold than the first level, and the last two levels will be as described above.

The above somewhat vague considerations do not specify the exact pattern selection parameters that should be used for each pass, especially the first three passes. These were only chosen after much trial and error, which would take too long to describe here. We do not have any theoretical justification for these parameters; they just seem to work well.

The table below shows the parameters used to generate the current set of TEX82 patterns, and the results obtained. For levels 2 and 4, the numbers in the "hyphens"


Table 5. Current TEX82 patterns.
column show the number of gnod and bad hyphens inhibited, respectively. The numbers in parentheses indicate the maximum length of patterns chosen at that level.

A total of 4919 patterns (actually only 4447 because some patterns appear more than once) were obtained, compiling into a suffix-compressed packed trie occupying 5943 locations, with 181 outputs. As shown in the table, the resulting algorithm finds $89.3 \%$ of the hyphens in the dictionary. This improves on the one and two level examples discussed above. The patterns were generated in 109 passes through the dictionary, requiring about 1 hour of CPU time.

## Examples

The complete list of hyphenation patterns currently used by TEX82 appears in the appendix. The digits appearing between the letters of a pattern indicate the hyphenation level, as discussed above.

Below we give some examples of the patterns in action. For each of the following words, we show the patterns that apply, the resulting hyphenation values, and the hyphenation obtained. Note that if more than one hyphenation value is specified for a given intercharacter position, then the higher value takes priority, in accordance with our level scheme. If the final value is odd, the position is an allowable hyphen point.

```
computer 4m1p pu2t 5pute put3er co4m5pu2t3er com-put-er
algorithm lig4 lgo3 1go 2ith 4hm alig4o3r2it4hm al-go-rithm
hyphenation hy3ph he2n hena4 hen5at ina n2at 1tio 2io
    hy3phe2n5a4t2ion hy-phen-ation
concatenation o2n on1c 1ca 1na n2at 1tio 2io
    co2n1cate1n2a1t2ion con-cate-na-tion
mathematics math3 ath5en th2e 1ma at1ic 4cs
    math5e1matil4cs math-e-mat-ics
```

```
typesetting type3 e1s2e 4t3t2 2t1in type3s2e4t3t2ing
    type-set-ting
program pr2 1gr pr2o1gram pro-gram
supercalifragilisticexpialidocious
    upe ric ica alli agii gil4 flif il4ist isiti st2i sitic
    1exp x3p pi3a 2i1a 12al 2id 1do 1ci 2io 2us
    su1per1cal1ifrag1il4is1t2iciex3p2i3al2i1do1c2io2us
    su-per-cal-ifrag-ilis-tic-ex-pi-ali-do-cious
```

- Below, we show a few interesting patterns. The reader may like to try figuring out what words they apply to. (The answers appear in the Appendix.)

| ain5o | hach4 | n3uin | 5spai |
| :--- | :--- | :--- | :--- |
| ay5al | h5elo | nyp4 | 4 tarc |
| ear5k | if4fr | o5a5les | 4 todo |
| e2mel | l5ogo | orew4 | uir4m |

And finally, the following patterns deserve mention:

```
3tex fon4t high5
```


## Chapter 5

## History and Conclusion

-The invention of the alphabet was one of the greatest advances in the history of civilization. However, the ancient Phoenicians probably did not anticipate the fact that, centuries later, the problem of word hyphenation would become a major headache for computer typesetters all over the world.

Most cultures have evolved a linear style of communication, whereby a train of thought is converted into a sequence of symbols, which are then laid out in neat rows on a page and shipped off to a laser printer.

The trouble was, as civilization progressed and words got longer and longer, it became occasionally necessary to split them across lines. At first hyphens were inserted at arbitrary places, but in order to avoid distracting breaks such as therapist, it was soon found preferable to divide words at syllable boundaries.

Modern practice is somewhat stricter, avoiding hyphenations that might cause the reader to pronounce a word incorrectly (e.g. considera-tion) or where a single letter is split from a component of a compound word (e.g. cardi-ovascular).

The first book on typesetting, Joseph Moxon's Mechanick Exercises (1683), mentions the need for hyphenation but does not give any rules for it. A few dictionaries had appeared by this time, but were usually just word lists. Eventually they began to show syllable divisions to aid in pronunciation, as well as hyphenation.

With the advent of computer typesetting, interest in the problem was renewed. Hyphenation is the 'H' of 'H \& J' (hyphenation and justification), which are the basic functions provided by any typesetting system. The need for automatic hyphenation presented a new and challenging problem to early systems designers.

Probably the first work on this problem, as well as many other aspects of computer typesetting, was done in the early 1050s by a French group led by G. D. Bafour. They developed a hyphenation algorithm for French, which was later adapted to English [U.S. Patent 2,762,485 (1955)].

Their method is quite simple. Ilyphenations are allowed anywhere in a word except among the following letter combinations: before two consonants, two vawela,
or $x$; between two vowels, consoriant-h, e-r, or s-s; after two consonants where the first is not $1, m, n, r$, or $s$; or after $c, j, q, v$, consonant- $w, m m, 1 r, n b, n f, n l, n m$, nn , or nr .

We tested this method on our pocket dictionary, and it found nearly 70 percent of the hyphens, but also about an equal amount of incorrect hyphens! Viewed in another way, about $65 \%$ of the erroneous hyphen positions are successfully inhibited, along with $30 \%$ of the correct hyphens. It turns out that a simple algorithm like this one works quite well in French; however for English this is not the case.

Other early work on automatic hyphenation is described in the proceedings of various conferences on computer typesetting (e.g. [30]). A good summary appears in [31], from which the quotes in the following paragraphs were taken.

At the Los Angeles Times, a sophisticated logical routine was developed based on the grammatical rules given in Webster's, carefully refined and adapted for computer implementation. Words were analyzed into vowel and consonant patterns which were classified into one of four types, and rules governing each type applied. Prefix, suffix, and other special case rules were also used. The results were reportedly "85-95 percent accurate", while the hyphenation logic occupies "only 5,000 positions of the 20,000 positions of the computer's magnetic core memory, less space than would be required to store 5008 -letter words averaging two hyphens per word."

Perry Publications in Florida developed a dictionary look-up method, along with their own dictionary. An in-core table mapped each word, depending on its first two letters, into a particular block of words on tape. For speed, the dictionary was divided between four tape units, and "since the RCA 301 can search tape in both directions," each tape drive maintained a "homing position" at the middle of the tape, with the most frequently searched blocks placed closest to the homing positions.

In addition, they observed that many words could be hyphenated after the 3rd, 5 th, or 7 th letters. So they removed all such words from the dictionary (saving some space), and if a word was not found in the dictionary, it was hyphenated after the 3rd, 5th, or 7th letter.

A hybrid approach was developed at the Oklahoma Publishing Company. First some logical analysis was used to determine the number of syllables, fnd to check if certain suffix and special case rules could be applied. Next the probability of hyphenation at each position in the word was estimated using three probability tables, and the most probable breakpoints were identified. (This seems to be the origin of the Time magazine algorithm described in Chapter 1.) An exception
dietionary handles the remaining cases; however there was some difference of opinion as to the size of the dictionary required to obtain satisfactory results.

Many other projects to develop hyphenation algorithms have remained proprietary or were never published. For example, IBM alone worked on "over 35 approaches to the simple problem of grammatical word division and hyphenation".

By now, we might have hoped that an "industry standard" hyphenation algorithm would exist. Indeed Berg's survey of computerized typesetting [32] contains a description of what could be considered a "generic" rule-based hyphenation algorithm (he doesn't say where it comes from). However, we have seen that any logical rountine must stop short of complete hyphenation, because of the generally illogical basis of English word division.

The trend in modern systems has been toward the hybrid approach, where a logical routine is supplemented by an extensive exception dictionary. Thus the incore algorithm serves to reduce the size of the dictionary, as well as the frequency of accessing it, as much as possible.

A number of hyphenation algorithms have also appeared in the computer science literature. A very simpie algorithm is described by Rich and Stone [33]. The two parts of the word must include a vowel, not counting a final e, es or ed. The new line cannot begin with a vowel or double consonant. No break is made between the letter pairs sh, gh, p, ch, th, wh, gr, pr, cr, tr, wr, br, fr, dr, vowel-r, vowel-n, or om. On our pocket dictionary, this method found about $70 \%$ of the hyphens with $45 \%$ error.

The algorithm used in the Bell Labs document compiler Roff is described by Wagner [34]. It uses suffix stripping, followed by digram analysis carried out in a back to front manner. In addition a more complicated scheme is described using four classes of digrams combined with an attempt to identify accented and nonaccented syllables, but this seemed to introduce too many errors. A version of the algorithm is described in [35]; interestingly, this reference uses the terms "hyphenating pattern" (referring to a Snobol string-matching pattern) as well as "inhibiting suffix".

Ocker [36], in a master's thesis, describes another algorithm based on the rules in Webster's dictionary. It includes recognition of prefixes, suffixes, and special letter combinations that help in determining accentuation, followed by an analysis of the "liquidity" of letter pairs to find the character pair corresponding to the greatest interruption of spoken sound.

Moitra et al [37] use an exception table, prefixes, suffixes, and a probabilistic break-value table. In addition they extend the usual notion of affixes to any letter
pattern that helps in hyphenation, including 'root words' (e.g. line, pot) intended to handle compound words.

## Patterns as paradigm

Our pattern matching approach to hyphenation is interesting for a number of reasons. It has proved to be very effective and also very appropriate for the problem. In addition, since the patterns are generated from the dictionary, it is easy to accommodate changes to the word list, as our hyphenation preferences change or as new words are added. More significantly, the pattern scheme can be readily applied to different languages, if we have a hyphenated word list for the language.

The effectiveness of pattern matching suggests that this paradigm may be useful in other applications as well. Indeed more general pattern matching systems and the related notions of production systems and augmented transition networks (ATN's) are often used in artificial intelligence applications, especially natural language processing. While AI programs try to understand sentences by analyzing word patterns, we try to hyphenate words by analyzing letter patterns.

One simple extension of patterns that we have not considered is the idea of character groups such as vowels and consonants, as used by nearly all other algorithmic approaches to hyphenation. This may seem like a serious omission, because a potentially useful meta-pattern like 'vowel-consonant-consonant-vowel' would then expand to $6 \times 20 \times 20 \times 6=14400$ patterns. However, it turns out that a suffixcompressed trie will reduce this to just $6+20+20+6=52$ trie nodes. So our methods can take some advantage of such "meta-patterns".

In addition, the use of inhibiting as well as hyphenating patterns seems quite powerful. These can be thought of as rules and exceptions, which is another common AI paradigm.

Concerning related work in AI, we must especially mention the Meta-DENDRAL program [38], which is designed to infer automatically rules for mass-spectrometry. An example of such a rule is $\mathrm{N}-\mathrm{C}-\mathrm{C}-\mathrm{C} \rightarrow \mathrm{N}-\mathrm{C} * \mathrm{C}-\mathrm{C}$, which says that if the molecular substructure on the left side is present, then a bond fragmentation may occur as indicated on the right side. Meta-DENDRAL analyzes a set of mass-spectral data points and tries to infer a set of fragmentation rules that can correctly predict the spectra of new molecules. The inference process starts with some fairly general rules and then refines them as necessary, using the experimental data as positive or negative evidence for the correctness of a rule.

The fragmentation rules can in general be considerably more complicated than our simple pattern rules for hyphenation. The molecular "pattern" can be a treelike or even cyclic structure, and there may be multiple fragmentations, possibly involving "migration" of a few atoms from one fragment to another. Furthermore, there are usually extra constraints on the form of rules, both to constrain the search and to make it more likely that meaningful or "interesting" rules will be generated. Still, there are some striking similarities between these ideas and our pattern-matching approach to hyphenation.

## Packed tries

Finally, the idea of packed tries deserves further investigation. An indexed trie can be viewed as a finite-state machine, where state transitions are performed by address calculation based on the current state and input character. This is extremely fast on most computers.

However indexing usually incurs'a substantial space penalty because of space reserved for pointers that are not used. Our packing technique, using the idea of storing the index character to distinguish transitions belonging to different states, combines the best features of both the linked and indexed representations, namely space and speed. We believe this is a fundamental idea.

There are various issues to be explored here. Some analysis of different packing methods would be interesting, especially for the handling of dynamic updates to a packed trie.

Our hyphenation trie extends a finite-state nachine with its hyphenation "actions". It would be interesting to consider other applications that can be handled by extending the basic finite-state framework, while maintaining as much of its speed as possible.

Another possibly interesting question concerns the size of the character and pointer fields in trie transitions. In our hyphenation trie half of the space is occupied by the pointers, while in our spelling checking examples from one-half to threefourths of the space is used for pointers, depending on the size of the dictionary. In the latter case it might be better to use a larger "character" size in the trie, in order to get a better balance between pointers and data.

When performing a search in a packed trie, following links will likely make us jump around in the trie in a somewhat random manner. This can be a disadvantage, both because of the need for large pointers, and also because of the lack of locality, which could degrade performance in a virtual memory environment. There are probably ways to inprove on this. For example, Fredkin [10] proposes an interesting ' $n$-dimensional binary trie' idea for reducing pointer size.

We have presented packed tries as a solution to the set representation problem, with special emphasis on data compression. It would be interesting to compare our results with other compression techniques, such as Huffman coding. Also, perhaps one could estimate the amount of information present in a hyphenated word list, as a lower bound on the size of any hyphenation algorithm.

Finally, our view of finite-state machines has been based on the underlying assumption of a computer with random-access memory. Addressing by indexing seems to provide power not available in some other models of computation, such as pointer machine, or comparisor-based models. On the other hand, a 'VLSI' or other hardware model (such as programmed logic arrays) can provide even greater power, eliminating the need for our perhaps contrived packing technique. But then other communication issues will be raised.

> If all problems of hyphenation have not been solved, at least some progress has been made since that night, when accordins to legend, an RCA Marketing Manager recelved a phone call from a disturbed customer. His 301 had Just hyphenated "God".
> - Paul E. Justus (1972)

TEX82 hyphenation patterns

| .ach4 | .ensa | -mo3ro | .under5 | ageso | a2n | apocs | asitr | aviter |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| .ad4dor | .eq5ui5t | , mu5ta | ,unie | 4agen | an3age | ap5ola | ur5a | av3ig |
| , af1t | .er4ri | .muta5b | .un5) | ag11 | Sanaly | apor81 | 2t | av50c |
| .a13t | .es3 | .n14c | .un50 | 4 ag 41 | a3nar | apos3t | t.3abl | afvor |
| .am5at | . 033 | . od2 | .un3u | agin | an3are | aps5es | at5ac | Samay |
| .an5c | .eye5 | . odd5 | .ap3 | a2go | anar4i | a3pu | at3alo | aw31 |
| .ang4 | . 1es3 | . 015 te | ,ure3 | 3agog | a3nati | aque5 | atsap | araly |
| .anism | .for5mer | .or5ato | .us5a | ag3oni | 4 and | 2a2r | ate5c | avst |
| .ants | .ga2 | .or3c | .ven4de | a5guer | andefs | ar3act | at5ech | axile |
| .an3te | .ge2 | .or1d | .ve5ra | ag5ul | an3dis | a5rade | at3ego | ax41d |
| .anti5s | . gen3t4 | .or3t | .*il5i | a4gy | an1dl | ar5adie | at3en. | ay5al |
| .ar5s | -ge5og | . 033 | . Yei | a3ha | an4dow | ar3al | at3era | aye4 |
| .arttio | . $\mathrm{g}^{15 \mathrm{a}}$ | .osttl | 4 ab . | ahe | a5nee | a5ramete | ater5a | ayst |
| .ar4ty | .g14b | .oth3 | a5bal | ah41 | a3nen | arandg | a5terna | az14er |
| .as3c | .go4r | .out3 | a 5 ban | a3ho | an5est. | ara3p | at3est | az51 |
| .as1p | .hand51 | .pud5al | abe2 | 12 | a3neu | ar4at | at5ev | 5 ba |
| ,as1s | .han5k | .pe5te | ab5erd | a5ia | 2ang | a5ratio | 4ath | bad5ger |
| .aster5 | .he2 | .pe5tit | abi5a | a3ic. | ang5ie | ar5ativ | ath5en | batge |
| .atom5 | .hero51 | .pi4e | ab5it5ab | a1517 | anigl | a5rau | a5then | balla |
| .au1d | .hen3 | .pio5n | ab5lat | 2414n | annilc | ar5avi | ataho | ban5dag |
| .av4i | .het3 | .pi2t | ab505112 | ain5in | a3nies | arami | ath5om | nfe |
| .awn4 | .hi3b | .pre3m | 4 abr | ain5o | an3131 | arbal4 | 4ati. | an3i |
| .ba4g | .hi3er | .ra4c | ab5rog | ait5en | an4ime | arfchan | a5tia | barbis |
| .ba5na | .hon5ey | .rantt | ab3ul | a1j | a5nimi | ar5dine | at5i5b | bari4a |
| .baste | .hon3o | .ratio5na | a4car | akion | a5nine | ar4dr | atiic | basfei |
| .ber4 | .hov5 | .ree2 | ac5ard | al5ab | an3io | ar5eas | at3if | lbat |
| . be5ra | .1d41 | .re5mit | ac5aro | al3ad | a3nip | a3ree | ation5ar | batz |
| . be3an | .idols | .res2 | a 5 ceou | a4lar | an31sh | ar3ent | atsitu | 2b1b |
| .be5sto | . im3m | .re5stat | acier | 4aldi | an31t | a5ress | aftog | b2be |
| .bri2 | .1m5pin | .ri4g | a5chet | 2ale | a3niu | ar4ti | a2tom | 3ber |
| .butfti | .in1 | .rit5u | 4 a 2 ci | al3end | an4kl1 | arifl | at5omiz | bi4na |
| .cam4pe | . 1 n 3 cl | .ro4q | a3cie | a4lenti | 5 anniz | aril | aftop | 4bld |
| .can5c | .ine2 | .ros5t | aciln | a5le5o | anos | ar5ial | a4tos | 4 b |
| .capa5b | .in2k | .row5d | a3cio | alli | an5ot | ar3ian | altr | eak4 |
| .car5ol | .in3s | .ru4d | ac5rob | al4ia. | anoth5 | a3riet | at5rop | beats |
| .ca4t | . $1 \times 5 \mathrm{r}$ | .sci3e | act5if | all4e | an2sa | r11] | at4ak | 4be2d |
| .ce4la | .1841 | .self5 | ac3ul | alslev | an4sco | ar5inat | at4tag | be3da |
| ,ch4 | . ju3r | .8ell5 | ac4um | 4 allic | an4sn | ar3io | at5te | be3de |
| .chil15i | .1a4cy | . 5 e2n | a2d | 4 alm | an2sp | ar2iz | atith | be3di |
| .ci2 | .1a4m | . se5rio | ad4din | a5log. | ans3po | ar2mi | a2tu | beigi |
| .cit5r | . 1 at5er | . 8 h2 | ad5er. | a417. | an4st | ar505d | at5ua | be5gu |
| .co3e | . lath5 | . 812 | 2adi | 4alys | an4sur | a5roni | at5ue | Ibel |
| .codr | . 102 | . 5 ing 4 | a3dia | 5a5lyst | antal4 | a3roo | at3ul | belli |
| .cos 5ner | . leg5e | . 8 t4 | ad3ica | 5alyt. | an4tie | ar2p | at3ura | -310 |
| .defmoi | . 1 en4 | .8ta5bl | adi4er | Salyz | 4 anto | ar3q | a2ty | $4 \mathrm{be5m}$ |
| .de3o | . lep5 | . $8 y^{2}$ | a3dio | 4ama | an2tr | arre4 | autb | be5nig |
| .de3ra | .lev1 | .ta4 | a3dit | am5ab | an4t\% | ar4sa | augh3 | be5nu |
| . de3ri | .114g | .te2 | a5diu | am3ag | an3ua | ar2sh | au3gu | 4best |
| . destc | .lig5a | .ten5an | ad4le | ama5ra | an3ul | 4as. | au412 | be3ap |
| .dictios | .112n | .th2 | ad3cre | am5asc | a5nur | astab | aun5d | be5str |
| .do4t | . 1130 | .t12 | ad5ran | a4matis | 430 | as3ant | au3r | 3bet |
| .du4c | .114t | .til4 | ad4su | 24m5ato | apar4 | ash14 | au5sib | bet5iz |
| . dumbs | .mag5a5 | .tim505 | 4adu | am5era | ap5at | a5sia. | aut5en | be5tr |
| .earth5 | .mal5o | .ting 4 | a3duc | an3ic | ap5ero | a3sib | uuith | be3tw |
| .eas31 | .man5a | .tin5k | ad5um | am5is | a3pher | a3sic | a2va | be3\% |
| .eb4 | .mar5ti | .ton4a | ae4r | am5ily | 4aphi | 5a5si4t | av3ag | be5yo |
| .eer 4 | .me2 | .to4p | aeri40 | am1in | a4pilla | ask3i: | a5van | 2 bl |
| .eg2 | .mer3c | .top5i | a2f | ami4no | ap5illar | as41 | ave4no. | 4b3h |
| .e15d | .me5ter | .tou5s | a114 | a2mo | ap3in | a4soc | av3era | b12b |
| .el3em | .mis1 | .trib5ut | a4gab | a5mon | ap3ita | as5ph | av5era | bi4d |
| .enam3 | .mist5i | .unla | aga4n | amor 51 | a3pitu | as4sh | av5ery | 3bio |
| .en3g | .mon3e | .un3ce | ag5ell | amp5en | a2pl | as3ten | av11 | bi5en |
| $\checkmark$ | $\checkmark$ | $\checkmark$ | - | , | , | , |  | , |


| b14er | b5uto <br> b1v | 3cheni | co3pa | 4dai | d2gy dih2 | Sdren dri4b | efben | elil4 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 b 11 | 4b5v | ch3er. | coupl | da2m2 | 5 di . | drill | e3br | ef5i5nite |
| b1311z | 5by. | ch3ers | 4 corb | dan 3 g | ! d413a | dro4p | efcad | Sefit |
| bina5ra | bys4 | Achin | coro3n | dard5 | dia5b | 4 drov | ecan5c | efor5es |
| bin4d | 1 ca | Schine. | coste | dark5 | difcam | Sdrupli | ecca5 | effuse. |
| bi5net | cab3in | ch5iness | covi | 4dary | duice | 4 dry | -1ce | 4egal |
| b13ogr | caibl | Schini | cove4 | 3dat | 3dict | 2d1:2 | ec5essa | eger4 |
| bi5ou | cach4 | Schio | cow5a | 4dativ | 3did | ds4p | ec21 | -g5ib |
| b12t | "ca5den | 3chit | coz5e | 4dato | 5di3en | disw | efcib | egife |
| 3bi3tio | 4 cag 4 | ch12z | co5zi | 5 dav 4 | d11f | d4ey | ec5ificat | egsing |
| bi3tr | 2c5ah | 3cho2 | c1q | dav5. | di3g. | d2th | ec5ifie | e5git5 |
| 3bit5wa | ca3lat | ch4ti | cras5t | 5day | di4lato | 1 du | ec5ify | -g 5 n |
| b5itz | calla | 1ci | 5crat. | dib | din | diula | ec3im | etgo. |
| b1j | callsia | Scia | Scratic | d5c | 1dina | du2c | eci4t | - Agos |
| bk4 | 4calo | ci2a5b | cre3at | d1d4 | 3dine. | diuca | e5cite | eginl |
| b212 | can5d | ciabr | 5 cred | 2 de . | 5 dini | duc5er | efclan | -5gar |
| blath5 | canfe | cisc | 4c3reta | deapt | di5niz | 4duct. | efclus | 50 gr |
| b4le. | canic | 4 cier | crest | deb5it | 1 dio | 4 diucts | e2col | elh4 |
| blent | can5is | Ecilic. | cri2 | defbon | dio5g | du5el | efcoma | eher4 |
| 5blesp | can3iz | 4cii | crist | decan 4 | di4pl | du4g | e4compe | 012 |
| t3lis | can4ty | ci4la | carin | do4cil | dir2 | d3ule | efcone | -5ic |
| b410 | cany4 | scili | crisf | de5com | dilre | dumibe | -2cor | -15d |
| blun4t | casper | 2cis | 5 criti | 281ed | dirt51 | du4n | ec3ora | elg2 |
| 4bis | car5om | 2 cin | croupl | 4 deo. | dis! | 4dup | eco5ro | -15gl |
| 4b3n | cast5er | clina | crop50 | de5il | 5disi | du4pe | elcr | -318b |
| bne5g | cas5tig | 3cinat | crosie | delise | dsis3t | d1v | efcrem | e3int |
| 3bod | 4casy | cin3en | cru4d | dol515q | d2iti | div | ectian | eling |
| bod3i | ca4th | cling | 4c3s2 | de510 | 1di1v | d2y | ecste | -5inst |
| bo4e | 4cativ | c5ing. | 2c1t | d4cm | d11 | 5dyn | elca | eir4d |
| bol3ic | cavEal | Scino | cta4b | 5 dea. | d5k2 | dy4ze | -4cul | -1t3. |
| bombl | c3e | cion4 | ct5ang | 3 demic | 4 d 51 a | dy $=5 \mathrm{p}$ | ec3ula | -13th |
| bonda | ccha5 | 4cipe | c5tant | dem51c. | 3d10. | e1a4b | 2 e 2 da | -5ity |
| bon5at | cci4a | ci3ph | c2to | de5all | 3dlod | -3act | 4ed3d | -1) |
| 3boo | ccompa5 | 4cipic | c3ter | defmons | 3dles. | -ad1 | -4dier | -4jud |
| Ebor. | ccon4 | 4cista | citicu | demor5 | 4dless | ead51. | edefs | ejbudi |
| 4 blora | ccoust | 4cisti | ctim3i | 1 den | 2 d 310 | -a4ge | 4edi | eki4n |
| borsd | 2 ce . | 2c1it | ctu4r | defnar | 4d51u | ea5ger | e3dia | ek4la |
| Sbore | 4 ced . | citsiz | c4tw | de3no | 2d1y | -241 | ed3ib | c1a |
| 5buri | 4cedea | Sciz | cud5 | denti5t | dis | ealser | - ed3ica | 41 c |
| $8 t .084$ | Scel | ck1 | caut | de3nu | 4din4 | eal3on | edsía | cllac |
| bsota | 8 cel . | ck31 | c4ui | delp | 1 do | eam3er | -dilt | elandd |
| boths | Scell | 1 c 414 | cu5ity | de3pa | 3do. | -5and | edi5z | -15ativ |
| botto | 1cen | Aclar | Sculi | depi4 | do5de | ear3a | 4edo | -41av |
| bounds | Scene | c5laratio | culutia | de2pu | 5 doe | -ar4c | -4dol | elaxat |
| 4 bp | $2 \cos 40$ | Eclare | Scultu | d3eq | 2 d 501 | ear5es | -don2 | e31ea |
| 4 brit | 4 ceni | clefm | cu2ma | d4erh | d40g | eartic | e4dri | -15ebra |
| broths | Scont | 4clic | c3ume | 5 dera | do4la | ear4il | -4dul | Selec |
| $2 \mathrm{bse2}$ | 3 cep | clias | cu4mi | dern5iz | dol14 | ear5k | ed5ulo | e4led |
| beort | co5ram | cly 4 | 3cun | der5: | do5lor | -ar2t | -02c | el3ega |
| 2 bt | 4 cesa | c5n | cu3pi | des2 | dom5iz | eart3. | eed3i | -51en |
| bt41 | 3 cessi | 1co | cu5py | d2es. | do3nat | ea5sp | -02f | e4l1ar |
| b4to | ces5a15b | co5ag | cur5a4b | delsc | don14 | -3ase | -0131 | -11es |
| bstr | $\operatorname{cosest}$ | coe2 | cu5ria | de2s50 | doo3d | cast3 | eefly | -128 |
| buf4fer | cetd | 2 cog | lcus | des3ti | dop4p | -a2t | -e2m | -121 |
| butga | c5efta | codgr | cuss41 | de3str | d4or | eat5en | -efna | -31ibe |
| bu311 | cert | cold | 3c4ut | de4su | 3dos | eath3i | eespl | -415ic. |
| bunis | 2ch | co3inc | cuitio | delt | 4d5out | e5atif | -02s4 | -13ica |
| buin | 4 ch . | colsi | $4 \mathrm{c5utiv}$ | de2to | dofy | c4a3tu | esat4 | -311er |
| bunt41 | $4 \mathrm{ch3ab}$ | Scolo | 4cuts | delv | 3 dox | ea27 | eefty | -15igib |
| bu3re | Schanic | col3or | Icy | dev3il | dip | eav3en | -5ex | *5118 |
| bus5io | ch5a5nia | com5er | cref | 4 dey | 1 dr | eav5i | e1f | -113ing |
| busato | che2 | conda | 1d2a | 4dif | drag50a | -av50 | eff3ere | -3110 |
| 5buet | cheap3 | ctore | Sda. | diga | <drai | 2elb | letf | e211. |
| 4 buta | Sclod | con 3 g | 2d3ab | d3gett | dred | efbel. | elific | e15ish |
| 3batio | clesio | con5t | dachs | dgli | drea5r | etbels | Sofici | e3liv3 |
|  | $\checkmark$ | $\checkmark$ | , | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |


| 4ella | e3ny. | er3ine | 4es2to | 1 fa | 111 n 4 | 4geno | go3ni | 1head |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ellab | 4en3z | elrio | e3ston | fa3bl | 1103re | Ageny | 5goo | 3hear |
| ellot | e5of | 4erit | 2estr | fab3r | 121 y 5 | 1 geo | go5riz | he4can |
| -5loc | 002g | er4iu | e5stro | fatce | 4 fm | ge3om | gor5ou | h5ecat |
| elsog | e4014 | eri4v | estruc5 | 4fag | 41 n | g4ery | 5 gos . | h4ed |
| el3op. | e3ol | efriva | e2sur | fain4 | 110 | 5gesi | govi | he5do5 |
| -12sh | eop3ar | er3m4 | es5urt | fallse | 5fon | geth5 | g3p | he3141 |
| el4ta | elor | erania | estir | 4fa4ma | fon4de | 4 geto | 1 gr | hel4lis |
| e5lud | eo3re | 4ernit | eta4b | fam5is | fondt | ge4ty | 4 grada | heluly |
| el5ug | eo5rol | Serniz | eten4d | 5far | forr | gefr | g4rai | h5elo |
| e4mac | eos4 | er3no | e3teo | far5th | forrat | 4g1g2 | gran2 | heasp |
| e4nag | e4ot | 2 ero | ethod3 | fa3ta | for5ay | g2ge | 5graph. | he 2 n |
| e5man | eo4to | er5ob | etiic | fa3the | fore5t | g3ger | g5rapher | henas |
| em5ana | e5out | -5roc | e5tide | 4fato | for4i | gglu5 | Egraphic | hen5at |
| em5b | -50\% | ero4r | etin4 | fault5 | fort5a | ggo4 | 4 graphy | heobr |
| -1me | e2pa | eriou | eti4no | 415b | fos5 | gh3in | 4 gray | hep5 |
| e2mel | e3pai | erls | e5tir | 41d | 415p | gh5out | grean | h4era |
| e4met | ep5anc | er3set | e5titio | 41 e . | 1ra4t | gh4to | 4 gress . | hera3p |
| em3ica | e5pel | ert3er | et5itiv | feas4 | 15rea | Egi. | 4 grit | her4ba |
| emite | e3pent | 4 ertl | 4etn | feath3 | fres5c | 1gi4a | $g 4 \mathrm{ro}$ | here5a |
| em5igra | ep5etitio | er3tw | et5ona | 104b | $1 \mathrm{fl2}$ | gia 5 r | grui4 | h3ern |
| em1in2 | ephes. | 4eru | e3tra | 4feca | $1 \mathrm{fil4}$ | glic | gs 2 | h5erax |
| em5ine | e4pli | eru4t | e3tre | 5 fect | frol5 | Sgicia | g5sto | h3ery |
| em3i3ni | elpo | 5 ¢rwau | et3ric | 2 fed | 213 s | g4ico | gth3 | hles |
| e4mis | e4prec | e1s4a | et5rif | fe3li | 21 t | gien5 | guam | he2s5p |
| em5ish | ep5reca | e4sage. | et3rog | fetmo | 14 to | 5gies. | 3 guard | heft |
| e5mise | e4pred | e4sages | et5ros | fen2d | 12ty | g114 | 2 gue | hetfed |
| em3iz | ep3reh | -82c | et3ua | fend5e | 31 u | g3imen | $5 \mathrm{gui5t}$ | hou4 |
| Seraniz | e3pro | e2sca | et5ya | ferl | fu5el | 3 g 4 in . | 3 gun | hif |
| emo4g | e4prob | es5can | et5z | 5ierr | fiug | gin5ge | 3 gus | h1h |
| emoni5o | ep4sh | e3scr | 4eu | fev4 | fu4min | 5g4ins | 4 gust | hisan |
| em3pi | ep5tisb | -85cu | e5un | 4111 | fu5ne | 5 gio | $\mathrm{g}^{31}$ | hi4co |
| efmul | e4put | -1820 | e3up | 14198 | fu3ri | 3 gir | 1 gy | hight |
| em5ula | ep5uta | -2eec | -u3ro | 14i1e | fusi4 | gir41 | 2 g 5 y 3 n | h4112 |
| omu3n | -1q | es5ecr | eus 4 | f5ifo. | fus4s | g3isl | gy5ra | himers |
| -3my | equi31 | es5enc | eute4 | 1295is | 4futa | g140 | h3ab41 | hiina |
| en5amo | e4q3uiss | efsert. | euti51 | 14117 | 117 | 5 giv | hach4 | hionde |
| -4nant | eria | e4serts | eu5tr | 12 fy | 1 ga | 3 giz | hae4m | h14p |
| ench4er | era4b | efserva | eva2p5 | 4fh | gat4 | g12 | hae4t | hir41 |
| en3dic | 4erand | 4esh | e2vas | 111 | 5 gal . | gla | h5agu | hi3ro |
| e5nea | or3ar | e3sha | ev5ast | 113a | 3 gall | glad5i | ha3la | hir ${ }^{\text {p }}$ |
| -5nee | 4erati. | esh5en | e5vea | 2i3ic. | ga31。 | 5 glas | hala3m | hir4r |
| en3en | 2 erb | -1si | ev3ell | 4i3ical | 2 gam | 1gle | hain | his3el |
| en5aro | erabl | e2sic | evel3o | 13ican | ga5met | g114b | hansci | his4s |
| en5esi | er3ch | -2sid | e5veng | Aificate | g5amo | g31ig | hanicy | hith5er |
| en5eat | or 4 che | es5iden | even4i | P3icen | gan51: | 3 glo | 5 hand. | hi2v |
| en3etr | 2 ers . | es5igna | evier | 113cer | ga3niz | glo3r | han4g | 4hk |
| e3new | e3real | e2s5im | e5verb | fic4i | gani5za | g1a | hang5er | 4h114 |
| en5ice | ere5co | esilin | e1vi | bilcia | 4gano | g4my | hang50 | hlant |
| e5nie | ere3in | esisute | ev31d | 5ifcie | gar5n4 | gn4a | h5a5niz | h210 |
| e5nil | er5el. | elitu | evi4l | 4fica | gasa4 | g4na. | han4k | hlo3ri |
| -3nio | er3eao | e5skin | e4vin | 113 cu | gath3 | gnetit | hanste | 4him |
| en3ish | er5ena | es4mi | evi4y | 115del | 4 gativ | g1n1 | hap31 | hmet4 |
| on3it | er5ence | -2801 | e5voc | 1ight5 | 4 gaz | g2nin | hap5t | 2hin |
| *5iu | 46rene | es3olv | -5vu | 11151 | $\mathrm{g}^{3 \mathrm{~b}}$ | g4nio | ha3ran | h5odis |
| Seniz | or3ent | e2son | -1va | 11115in | gd4 | gino | ha5ras | h5ods |
| 4 enn | ere4q | es5ona | elvag | 4ifly | 2 ge . | g4non | har2d | ho4g |
| 4eno | or5ess | elep | e5vee | 2 fin | 2 ged | 1 go | hard3e | hoge4 |
| -no4g | ar3est | es3per | e3wh | Sifina | geez4 | 3 go . | har4le | hol5ar |
| -4nos | -ret4 | es5pira | emil5 | 1 in 2 d 5 | gelin | gob5 | harp5en | 3hol4e |
| en3ov | orlh | es4pre | er3ing | 112ne | ge5'is | 5 goe | har5ter | ho4ma |
| entsw | erli | 2ess | e3wit | 11 in 3 g | ge51iz | 3 g 404 g | has5s | home3 |
| ent5age | eiria4 | -84si4b | 1exp | $1 \mathrm{in4n}$ | 4 gely | 503is | haun4 | hon4a |
| 4enthes | Serick | estan4 | Seyc | fis4ti | 1 gen | gon2 | 5 haz | ho5ny |
| en3ua | e3rien | es3tig | 5eye. | 1412 | getnat | 4 g 303 na | haz3a | 3heod |
| -n5ut | ori4er | ea5tis | eys4 | 151ess | ge5niz | gondo5 | hib | hoond |
| $\checkmark$ | , | $\checkmark$ | ) | $\checkmark$ | $\checkmark$ | $\checkmark$ |  | $\checkmark$ |


| hor5at | 4iceo | ig3in | 4ingu | irfmin | it3uat | k1m | 31ess | 13108 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ho5ris | 4ich | ig3it | 2 ini | 1r04g | 15 tud | k5nes | 51 ess . | 131el |
| hort30 | 21 ci | 14 g 41 | 15 ni . | 5iron. | it3ul | 1kÊno | 13eva | 131efn |
| ho5ru | 15 cid | 12go | 14nia | 1 F 5 ul | 4itz. | ko5r | levier. | 13104t |
| hos4e | ic5ina | ig3or | in3io | 218. | 110 | kosh4 | levAera | 1121 |
| ho5sen | 12cip | ig5ot | in11s | 185ag | 2iv | k3ou | leviers | 121in4 |
| hosip | ic3ipa | 15 gre | i5nite. | 1s3ar | iv3ell | kro5n | 31 ey | 151ina |
| 1hous. | 14 cly | igu5i | 5 initio | 1sas5 | iv3en. | 4k1s2 | 41eye | 1140 |
| house3 | 12c5oc | igivr | in3ity | 2isic | 14v3er. | k4sc | 212 | $110 q u 15$ |
| hov5el | 411cr | 13 h | 4ink | 1 s 3 ch | 14vers. | ks 41 | 15 fr | 1150ut |
| 4h5p | Eicra | 41514 | 4inl | 4 fige | iv511. | $\mathrm{k4} \mathrm{sy}$ | 411g4 | 1510\% |
| 4 hr 4 | 14 cry | 13 j | 2 nn | is3er | 1v5io | k5t | 15 ga | 21. |
| hreo5 | 1 c 4 to | 4ik | 211no | 3181 | 1v1it | k17 | 1 gar 3 | 15mot |
| hro5niz | ictu2 | i11a | 14no4c | 1s5han | 15vore | lab3ic | 14ges | 1m3ing |
| hro3po | ic4t3ua | 113a4b | ino4s | is3hon | 17303ro | 14abo | lgo3 | 14mod |
| 4h1s2 | ic3ula | 141ade | 14 not | ish5op | 14v30t | laci4 | 213h | 1mont |
| h4gh | ictum | 1215 am | 2ins | 1s3ib | 415\% | 14ade | 114 ag | 211n2 |
| h4tar | 1c5uo | 11a5ra | - in3se | 1814 d | $1 \times 40$ | 1a3dy | 112am | 310. |
| htion | 13 cur | i3leg | insur5a | 15sis | 41y | lag4n | liar5iz | lob5al |
| ht5es | 2id | iller | 2 int. | is5itiv | 4izar | $1 \mathrm{am3o}$ | 114 as | 104 ci |
| h4ty | 14dai | ilev4 | 2 insth | 4is4k | 1214 | 31 and | 114ato | 4101 |
| hu4g | Id5anc | 1151 | indu | islan4 | 5izont | lau4dl | 115 bi | 3logic |
| hu4min | 1d5d | 1111 | 15nus | 4icme | 5 ja | lan5et | 5licio | 150go |
| hun5ke | ide3al | 1131a | 4 iny | 1250 | jac4q | lan4te | 114cor | 310 g |
| hun4t | ide4s | 112ib | 210 | 1605 mer | fa4p | lar4g | 4lics | 10m3er |
| hus3t4 | 12di | 11310 | 410. | 181p | 1 je | lar3i | 4lict. | 5long |
| hu4t | idsian | 114ist | 10go4 | 182pi | jer5s | lase | 141cu | 10 sif |
| h1\% | idi4ar | 2ilit | 102 gr | 154 py | 4jestie | la5tan | 131cy | 1303niz |
| h4mart | 15die | 112iz | 1101 | 4isis | 4jesty | 41 ateli | 131da | lood5 |
| hy3pe | 1d3io | 1115 ab | 104m | istsal | јеv3 | 4lativ | 11d5er | 510 pe . |
| hy3ph | 1di5 \% | 411n | ion3at | issent | jo4p | 4lav | 3lidi | lop31 |
| hy2s | idilt | 11309 | ionfery | istses | 5judg | 12474a | 1113er | 130pm |
| 211a | id5iu | 114 ty | ion3i | issta. | 3 ka . | 211b | 14119 | lorat |
| 12al | 13dle | 115ur | 105ph | isite | k3ab | lbin4 | 11411 | 104rato |
| iam4 | 14dom | 1137 | ior31 | isiti | k5ag | 411c2 | 5ligate | 105 ri ( |
| iam5ete | id3ow | 14 mag | 1408 | ist4ly | kais4 | lces | 311gh | lor5ox |
| 12 an | 14 dr | im3age | 105th | 4istral | kal4 | 13 ci | li4gra | 5108. |
| tianc | 12 du | ima5ry | 150ti | 12 su | k1b | 21d | 311k | 1035et |
| 1an31 | 1d5uo | imenta 5 r | io4to | 185us | k2ed | 12 de | 414141 | 5losophis |
| 4ian4t | 2104 | 4imet | 140ur | 4 ita . | 1kee | 1d4ere | . 11 mabl | 510aphy |
| 1a5pe | ied4e | 1m11 | 21p | ita4bi | ke4g | 1d4eri | 11.331 | 1084t |
| 1ass4 | 51 e ga | im5ida | ipe4 | 14 tag | ke511 | 1 d 14 | 114mo | 1osta |
| 14ativ | ield3 | imi5le | iphras4 | 4ita5m | k3en4d | 1d518 | 141m4p | loun5d |
| ia4tric | ien5a4 | 15 mini | ip3i | 13 tan | kler | 13 dr | lifina | 2lout |
| 14atu | ienfe | 4imit | 1 p 41 c | 13tat | kes4 | 14 dri | 114180 | 410 V |
| 1 bel | 15enn | 1 m 4 ni | ip4re4 | 2ite | k3est. | 1e2a | lin3ea | 21p |
| ib3era | 13 enti | 13 mon | 1p3ul | it3era | ke4ty | 104 bi | 11 n 31 | $1 \mathrm{pa5b}$ |
| ib5ert | 11er. | 12 mu | 13 qua | i5teri | k31 | left5 | link5er | 13 pha |
| ib5ia | 13esc | im3ula | iq5uef | it4es | kh4 | $510 g$. | 1150g | 15 phi |
| ib3in | 11est | 2in. | iq3uid | 21th | k11 | 51 lgg | 41419 | 1p5ing |
| ibsit. | 13et | 14 n 3 au | iq3ui3t | 11ti | 5ki. | letmat | 1184p | 13pit |
| ib5ite | 411. | 4inav | 41 r | 4itia | 5 k 2 ic | lem5atic | 111t | 14pl |
| 11bl | if5ero | incel4 | 11 ra | 412tic | k4111 | 41 en . | 121t. | 15pr |
| 1 b 311 | iff5en | in3cer | ira4b | it3ica | kilo5 | 31 enc | slitica | 411 r |
| 15 bo | 114ir | 4 ind | 14 rac | 5i5tick | k4is | 51 ene. | 1515tics | 211-2 |
| 11 br | 4ific. | in5dling | ird5e | it3ig | k4in. | 1lent | 1iv3er | 14:c |
| $12 \mathrm{b5ri}$ | 13110 | 2ino | iresde | 1 15ill | kin4de | 1 e 3 ph | 1118 | 1280 |
| 15bun | 1311 | 13nee | 14ref | 12tim | k5iness | le4pr | 415 | 14sie |
| Aican | 4ift | iner4ar | 14 rel 4 | 2 2tio | kin4g | lera5b | 1ka3 | 416 |
| Sicap | 2ig | 15ness | 14 res | 4itis | $k 14 p$ | lerfe | 13kal | 1t5ag |
| 4icar | iga5b | 4 inga | ir5gi | 14tism | kis4 | 31 rag . | 1ka4t | 1 tane5 |
| 14car. | 1 g 3 era | 4inge | $1{ }_{151}$ | 12t505m | k5ish | 314 eri | 111 | 11te |
| i4cara | $1 \mathrm{ght3i}$ | in5gen | 1ri5de | 4iton | kx4 | 14ero | 141av | 1 ten4 |
| 1cas5 | 4igi | 4ing 1 | ir4is | i4tram | $k 11$ | les2 | 121. | 1 toras |
| 14cay | 13gib | in5gling | iri3tu | it5ry | 4kley | 105sco | 151ea | 1th31 |
| iccul | ig3il | 4ingo | 5i5r2iz | 4itt | 4Kly | 51esq | 131*e | 15ties. |
| $\sqrt{ }$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |


| 1 1.84 | 4me. | manin | n5act | ne4po | nk3in | nti2f | -211 | 013ume |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| litr | 2 med | miso | nag5er. | $n \mathrm{n} 2 \mathrm{q}$ | nikl | n3tine | of5ite | ol3un |
| 1212 | 4 med . | 1 mo | nak4 | nier | 4 n 11 | n4t3ing | ofitat | o5lue |
| 1tur3a | 5 media | 4 mocr | nafli | nera5b | n5m | nt14p | 02 g 5 abr | -12v |
| 1u5a | me3die | 5 mocratiz | na51ia | n4erar | nme4 | ntrol511 | og5ativ | 0215 |
| 1 u 3 br | m5e5dy | mo2d1 | 4 nalt | n2ere | nmet4 | $n \mathrm{nta}$ | -4gato | Om5ah |
| luch4 | me2g | mo4go | na5mit | $n 4$ er51 | 4 n 1 n 2 | ntu3me | olge | oma51 |
| 143 ci | mel5on | mois2 | n2an | ner4r | nne4 | nula | -5gene | om5atiz |
| 103 en | mel4t | moi5se | nanci4 | ines | nni3al | nu4d | 05geo | om2be |
| 1 l 4 | me2m | 4mok | nan4it | 2nes. | $\mathrm{nni4v}$ | nu5en | 04 ger | om4bl |
| 1051d | mem103 | mo5lest | nank4 | 4nesp | nob4l | nuf4ie | o3gie | -2me |
| 104ma | 1 men | mo3me | nar3c | 2nest | no3ble | n3uin | 101gis | om3ena |
| 5lumi | men4a | mon5et | 4 nare | 4 nest | $n 50 \mathrm{cl}$ | 3nu3it | og3it | om5erse |
| 15 um . | men5ac | mon5ge | nar3i | 3 netic | 4 n 302 d | $n 4 \mathrm{um}$ | 04 gl | 04met |
| 5lumnia | men4de | moni3a | nar41 | neft | 3noe | nuime | -5g2ly | ombetry |
| lu30 | 4 mene | mon41sm | n5arm | n5eve | 4nog | n5umi | 30 gniz | -3mia |
| 1vo3r | men4i | mon4ist | n4as | ne4t | noge4 | $3 \mathrm{nu4n}$ | -4gro | om3ic. |
| 41 up | mens4 | mo3niz | nas4c | n31 | nois51 | n3uo | ogu5i | om3ica |
| luss 4 | mensu5 | monol4 | nas5ti | $n 4 \mathrm{gab}$ | no514i | nu3tr | 10 gy | 05aid |
| 1us3te | 3 ment | mo3ny. | n2at | n3gel | 5nologie | niv2 | 20 gyn | omiln |
| 11ut | menste | mo2r | na3tal | ngean4e | 3nomic | n1w4 | 01 h 2 | obaini |
| 15ven | me5on | 4mora. | nato5miz | n5gere | n505miz | nym4 | ohab5 | 5ommend |
| 15 vet 4 | m5eram | mos2 | n2au | $n 3 \mathrm{geri}$ | no4mo | nyp4 | 012 | omonge |
| 211\% | 2mes | mo5sey | nau3se | ng 5ha | no3my | 4 nz | oic3es | 04mon |
| 117 | 3mesti | mo3sp | 3 naut | n3gib | no4n | n3za | oi3der | om3pi |
| 41 ya | me4ta | moth3 | nav4e | ngin | non4ag | 402 | 01114 | ompro5 |
| 417 b | met3al | m5ouf | 4 n 1 b 4 | n5git | non51 | oad3 | oig4 | 02 n |
| 1y5me | meite | 3 mous | near5 | n4gla | n5oniz | 05a51es | oislet | onla |
| 173no | me5thi | mo2v | n4ces. | ngova | 4nop | oard3 | o3ing | on4ac |
| 2174 | metr | 4 m 1 p | n3cha | ng 5 sh | 5 nop5o5li | oas4e | oint5er | o3nan |
| 15 yse | 5 metric | mpara5 | n5cheo | nigu | nor5ab | oast5e | 051am | onic |
| 1 ma | me5trie | mpa5rab | $n 5 \mathrm{chil}$ | $n 4 \mathrm{gum}$ | no4rary | oat51 | $015 s$ on | 3 oncll |
| 2 mab | me3try | mpar5i | n3chis | n 2 gy | 4 nosc | ob3a3b | oist5en | 2ond |
| ma2ca | me4t | m3pet | nciln | $4 \mathrm{nih4}$ | nos4e | 05 bar | oi3ter | on5do |
| ma5chine | 4 mlf | mphas 4 | nctit | nhea | nos5t | obe41 | 05j | o3nen |
| ma4cl | 2mh | m2pi | neour5a | nhab3 | no5ta | 0161 | 20k | on5est |
| mag5in | 5mi. | mpi4a | nicr | The4 | inou | 02bin | o3ken | on4gu |
| 5 magn | mi3a | mp5ies | nicu | 3 n 41 a | 3noun | ob5ing | ok5ie | onlle |
| 2 mah | mid4a | m4piin | n4dai | ni3an | nov3el3 | 03 br | 011 a | $03 n 10$ |
| maid5 | mid4g | m5pir | n5dan | ni4ap | now13 | ob3ul | 041an | on11: |
| 4mald | mig4 | mp5is | nide | ni3ta | $n 1 p 4$ | olce | 0lass 4 | 05 niv |
| ma3lig | Smilia | mpo3ri | nd5est. | ni401 | npi4 | och4 | 012 d | on3key |
| ma5lin | m5i51io | mpos5ite | ndi4b | $n 14 \mathrm{~d}$ | npre4c | o3chet | oldie | on4odi |
| mal411 | m4111 | ${ }^{\text {m }}$ pous | n5d2it | n15di | n19 | ocif3 | o13er | on3ony |
| mal4ty | min4a | mpov5 | nidit | ni4er | nir | 04 cil | -31esc | on3s |
| Smania | 3mind | mpatr | n3diz | ni291 | nru4 | 04clam | 031et | onspi4 |
| man5is | m5inee | m2py | n5due | nisficat | 2n1s2 | 04cod | 01411 | onspir5a |
| man3iz | ming ${ }^{\text {d }}$ | 4 m 3 r | ndu4r | n 51 gr | ne5ab | oc3rac | 0121 | onsu4 |
| 4 map | min5gli | 4m1s2 | nd2we | nik4 | nsati4 | oc5ratiz | -311a | onten4 |
| masirine. | msingly | $\mathrm{m4sh}$ | 2ne. | $n 11 \mathrm{~m}$ | ns4c | ocre3 | -31ice | on3t4i |
| ma5riz | mindt | m5ai | n3ear | ni3miz | n2se | 5ocrit | ol5id. | ontif5 |
| mar4ly | m4inu | 4mt | ne2b | $n 1 i n$ | n4s3es | octor5a | -31141 | on5um |
| mar3y | miots | 1 mu | neb3u | 5 nine. | neidi | oc3ula | 05111 | onva5 |
| ma5sce | mis | mula5r4 | ne2c | nin4g | nsigd | o5cure | ol3ing | 002 |
| mas4e | mister. | 5 mult | 5 neck | $n 140$ | n2s1 | od5ded | 05110 | ood5e |
| masit | mis51 | multi3 | 2ned | $5 \mathrm{nis}$. | 3m | od3ic | 051is. | ood5i |
| 5 mate | misati | 3 mum | ne4gat | nis4ta | n4soc | odi3o | -13ish | 004k |
| math3 | m5istry | mun2 | neg5ativ | n2it | ns4pe | -2do4 | -51ite | oop3i |
| ma3tis | 4mith | 4 mup | Snege | n4ith | n5spi | odor3 | 0511tio | o3ord |
| 4matiza | u2iz | mu4u | ne4la | 3 nitio | nsta5bl | od5uct. | 051 l | oost5 |
| 4m1b | 4 mk | 4mi | nelsiz | n3itor | $n 1 t$ | od5ucts | olli4e | -2pa |
| mba4t5 | 4 mll | Ina | ne5mi | ni3tr | nta4b | 04 el | 0150 giz | ope5d |
| m5bil | min | 2n1a2b | ne4mo | $n 1 j$ | $n$ nter3s | 05eng | 0104 r | opler |
| m4b3ing | manary | n4abu | Inen | 4 nk 2 | nt21 | o3er | -15pl | 3 popera |
| mbidy | 4 min | 4 nac. | 4 nene | n5kero | n5tib | 0efta | -12t | 4operag |
| 4 mbe | mata/ | na4ca | 3neo | n3ket | nti4or | -3ev | ol3ub | 20ph |
| $V$ | $\sqrt{ }$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |


| -5phan | oftes | pear41 | pind4 | proit | rb40 | rev5ols | rivset | repeat |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -5pher | 40th | pe2c | p4ino | 2p1s3 | ric | re4wh | riv3i | rp5er. |
| op3ing | oth5esi | 2p2ed | 3pi10 | P2ss | r2ce | r18 | r3j | r3pet |
| -3pit | oth314 | 3 pede | pion4 | p 4 h | rcent rich | 5104 | r3ket | rp4h4 |
| ospon | ot3ic. | 3pedi | p3ith | p4sib | r3charrah | rify | rx4le | rp3ing |
| 04 posi | ot5ica | pediat | pistha | 2p1t | rchfor | rg2 | rkilia | r3po |
| 01 pr | o3tice | pedic | p12tu | pt5ab | racidb | rg3er | r11 | rir4 |
| opin | ostıf | p4ee | 2p3k2 | p2to | resit | r3get | rle4 | rrefe |
| opy5 | 03tis | peesd | 1 p 212 | p2th | rcuas | r3gic | r2led | rreti |
| 019 | -to5: | pek4 | 3plan | pt13m | r4dal | rgitn | ralig | r4reo |
| olra | OU2 | pefla | plas5t | ptufr | rd21 | rg3ing | r41s | rredst |
| obra. | ou3bl | pel14e | pl13a | p4t\% | rdi4a | r5gis | r15iah | r140 |
| 04r3ag | ouch51 | peinan | pliser | pubs | rdifor | r5git | r 3104 | rri4v |
| or5aliz | ou5et | pfenc | 4plig | puef | rdin4 | rigl | ria | rond |
| or5ange | ou41 | pen4th | pli4n | puit | rd3ing | rgodn | rase | rrast |
| or*5a | ounc5er | pe5on | p1014 | pul3e | 2 re . | r3gu | r2me | rrys4 |
| orreal | oun2d | pfera. | plutm | pu4m | relal | rh4 | 3men | 4 ra 2 |
| or3el | ou5V | pera5bl | plumab | pu2n | re3an | 4 rh . | m5ers | risa |
| ore5sh | ovien | p4erag | 4pim | pur4r | re5arr | 4 rhal | ra3ing | rastid |
| orbest. | over 4 ne | p4eri | 2 p 3 n | 5pus | 5 reay | ri3a | 4ining, | a4c |
| orewf | over3s | perisst | podc | pu2t | refay | riadb | 4盟10 | 280 |
| origu | ov4ert | per4mal | 5pod. | 5puter | r5ebrat | ri4ag | r3mit | Ssec |
| 405ria | o3vis | perme5 p4ern | po5em | put3er | rec5oll | ridb ribSa | ramy ranar | rsefcr |
| 05 r 11 | 05 T 401 | per3o | $5 \mathrm{po4g}$ | putited | refcre | ric5as | r3nel | Ses |
| oriln | ow3der | per3t1 | poin2 | putatin | 2 r 2 ed | raice | Aner | -05v2 |
| olrio | -073el | pe5ru | Spoint | p3\% | relde | 4rici | r5net | rioh |
| or3ity | -\%5est | periv | poly5t | qu2 | re3dis | 5ricid | r3ney | r5sha |
| 03rin | 0.11 | pe2t | po4ni | qua5v | red5it | ri4cie | r5nic | r1ai |
| or2mi | orn51 | pe5ton | po4p | $2 q u e$. | 4fac | raico | rinis4 | rasilb |
| orn2e | 0470 | pe5tiz | ipaor | 3 quer | re2fe | rid5er | r3nit | reons |
| 05 rof | oyla | 4 p ? | po4ry | 3 quet | re5fer, | ri3enc | r3niv | risp |
| or3oug | 1 pa | 4pg | 1pos | 2 rab | re3fi | ri3ent | rno4 | 5 sw |
| or5pe | pa4ca | 4 ph . | posis | ra3bi | resty | riler | rama | rtach |
| Sorrh | pa4ce | phar51 | prot | rach4e | reg3ia | r15et | r3nu | r4tag |
| or4se | pac4t | phe3no | po4ta | r 5acl | re5it | rig5an | rob31 | r3teb |
| ora5en | p4ad | ph4er | 5 poun | raf5i1 | relli | 5 rigi | r2oc | rten4d |
| orst4 | 5 pagan | ph4es. | 4p1p | ral4t | re5lu | r113iz | 03 cr | rte50 |
| or 3thi | p3agat | philc | ppa5ra | r2aj. | renta | 5riman | 040 | ret |
| or3thy | p4ai | 5phie | p2pe | ra410 | rente | rim5i | rolfe | rt5ib |
| or4ty | pain4 | ph5ing | psped | ram3et | reio | 3rimo | ro5ill | rtisd |
| o5rum | p4al | 5phisti | p5pel | r2ami | re5pin | rim4pe | ok2 | ritior |
| olry | panfa | 3 phiz | p3pen | rane50 | re4posi | r2ina | ro5ker | r3tig |
| os3al | pan3el | ph21 | p3per | rantge | relpu | 5rina. | 5role. | rtil31 |
| os2c | pan4ty | 3 phob | p3pet | raani | rier4 | rin4d | rom5eto | rtil41 |
| -stce | pa3ny | 3phone | ppo5site | ra5no | r4eri | rin4e | romil | ratily |
| 03scop | palp | 5phoni | pr2 | rap3er | ro4 | rin4g | romip | r4tist |
| 4oscopl | pa4pu | pho4r | pray4e | 3 raphy | 5ru | rilo | ronfal | ritiv |
| 05 ccr | para5bl | Aphs | 5preci | rar5c | es. | 5riph | ron4e | 3 tri |
| 084140 | par5age | ph3t | pre5co | rare4 | restipi | riph5e | 5 n 1 s | rtroph4 |
| 0s5itiv | par5di | Sphu | pre3em | rar5ef | ress5ib | ri2pl | ronfta | rtish |
| os31to | 3 Pare | 1 phy | preitac | 4 raril | 2t | rip5lic | 1 room | 3a |
| os3ity | par5el | pi3a | pre4la | r2as | re5stal | r4iq | 5 root | u3e41 |
| osifu | p4a4ri | pian4 | pre3r | ration4 | 3str | r21: | ro3pel | u3en |
| - 041 | par4is | pi4cio | p3rese | rau4t | refter | r4is. | rop3ic | ruAgl |
| 0280 | paite | pi4cy | 3 press | ra5vai | 4ti4z | riskc | ror3i | ru3in |
| 084 pa | pa5ter | p4id | pre5ten | rav3el | re3tri | r31sh | ro5ro | rum3pl |
| 0s4po | 5pathic | p5ida | pre3v | ra5zie | 12 | ris4p | ros5per | 2n |
| os2ta | pa5thy | pi3de | 5pri4e | r1b | re5uti | ri3ta3b | ros4s | unk5 |
| 05stati | pa4tric | Spidi | prin4t3 | r4hab | rev2 | r5ited. | ro4the | rundty |
| 0s5til | pav4 | 3 piec | pri4s | 14 bag | retval | rit5er. | rodty | r5usc |
| Os5tit | 3 pay | pi3en | pris3o | rb12 | revel | rit5ers | r04va | uti5a |
| 04 tan | 4p1b | pi4grap | p3roca | rbi4i | r5ev5er. | rit3ic | rov5el | ve |
| otele4g | pd4 | pi3lo | prof5it | r2bin | re5vers | ri2tu | rox5 | rvelut |
| ot3er. | 4pe. | pi2n | pro3l | r5bine | re5veri | rit5ur | r1p | r3ven |
| ot5ers | 3pe4a | piin. | pros3e | rb5ing. | re5vil | riv5el | r4pea | ry5er. |
| $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | J | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |


| r5vest | e5ened | 2*19 | 8tag | tal3i | 2 tIf | t510 | 4 tap4 | ugbia |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| r3vey | 80.55 | s3ma | 2tal | 4 talk | 4tig | 4t19 | $5 \mathrm{tu3i}$ | 2412 |
| r3vic | -5enin | small3 | $3 \operatorname{tam41}$ | tal4lis | 2th. | tmed | Stum | a11518 |
| rvi4y | 4sentd | sman3 | 5 stand | ta5log | than4 | 2 LIn 2 | unis | 14 n |
| r3vo | 4eentl | smel4 | $84 \mathrm{ta4}$ p | ta5mo | th2e | 1 to | 2t3up. | 1ing |
| r1\% | sep3a3 | 5 men | 5 stat. | $\boldsymbol{t a n} 4 \mathrm{do}$ | 4thea | to3b | Sture | $1{ }^{16}$ |
| ryde | 4sior. | 5smith | 4 ted | tanta 3 | th3eas | to5crat | 5 turi | uital |
| 5rynge | sterl | mmol5d4 | stern51 | ta5per | the5at | 4 todo | tur3is | 1v3 |
| ry3t | serso | -1n4 | e5tero | ta 5 pl | the3is | 2 tol | tur 50 | ivior. |
| sa2 | 4servo | 1so | ste2v | tar4a | 3thet | to2gr | tu5ry | v5) |
| 2alab | -1efs | eofce | stew5a | 4 tare | th5ic. | tosic | Stus | 40k |
| 5asack | se5sh | softs | 3the | 4 tare | th5ica | to2ma | 4tv | u11a |
| sac3ri | ees5t | so4lab | st2i | ta3riz | 4 thil | tom4b | t=4 | ala ${ }^{\text {b }}$ |
| s3act | $5 \mathrm{eo5um}$ | sol3d2 | stti. | 40 | Sthink | to3my | 4tiva | n5lati |
| 5 sai | 5sev | so3lic | Etia | ta5sy | 4thl | ton4all | twis 4 | alch4 |
| asarat | sev3en | 58017 | attic | Static | th5ode | tosnat | 4two | 5ulche |
| salfa | ser41 | 3som | 5stick | taftur | Sthodic | 4 tono | 1ty | ul3der |
| a 2510 | 5eex | 3s4on. | stito | taun4 | 4 thoo | 4tony | 4tya | u140 |
| cal4t | 4831 | cona4 | s3tif | tav4 | thor51t | to2ra | 2tyl | allen |
| Ssanc | 283g | son4g | stsing | 2tav | tho5riz | to3rie | types | v14gi |
| sansde | 2 h | 340p | Satir | tax4ia | 2 the | tor5iz | ty5ph | 4121 |
| slay | 2 hh . | Seophic | sitle | 2tib | 1tia | tos2 | 4 tz | u511a |
| sa5ta | shler | -5ophiz | 5stock | 4 tc | ti4ab | Stour | ta4e | ul3ing |
| 5astio | Esher | s5ophy | atom3a | t4ch | ti4ato | 4 tout | 4 uab | u1518h |
| eat3u | shisn | sor5c | 5stone | tch5et | 2ti2b | to3war | uact | ultlar |
| saut | oh3io | cor5d | stop | 4t1d | 4tick | 4t1p | ua5na | u14146 |
| sa5vor | 3ship | 4sov | 3 store | 4 to. | t4ico | 1 tra | uan4i | nl418s |
| 5amp | shiv5 | s0571 | stur | tead4i | tificiu | tra3b | uar5ant | $4 \mathrm{ul3}$ |
| 4.5b | shot | 2spa | s4trad | 4 teat | 5tidi | tra5ch | uar2d | 41140 |
| acandt5 | ch5old | 5 spal | Sstratu | teced | 3tion | traci4 | uar31 | 4uls |
| ceadp | shon3 | spain | s4tray | Stect | til2 | rac4it | uar3t | ula5es |
| acav5 | shor 4 | spen4d | sttrid | 2t1ed | ti5fy | ractto | u1at | uliti |
| siced | short5 | 2s5peo | 4 stry | te5di | 2tig | ras4 | uav4 | ultras |
| 4scel | 4sh\% | 2sper | 4st3\% | 1 teo | 5 5igu | ra5ven | ub4e | 4ultu |
| Stces | silb | 2phe | s2ty | teg 4 | tillsin | rav5es5 | ufbel | n310 |
| sch2 | sice | 3 3pher | 1 su | te5ger | 1tim | re5t | 3 ber | u15ul |
| Acho | 3aide. | ephos | su1al | te5gi | 4timp | e4m | lbero | 4157 |
| 3 scio | 5sides | spil4 | su4b3 | 3 tel. | tim5ul | trem51 | ulb4i | um5ab |
| Sscin4d | Esidi | sp5ing | su2g3 | teli4 | 2 t 1 in | 5 tria | ufb5ing | um4bi |
| acle5 | si5diz | 4 spio | su5is | 5 tels | t2ina | tri5ces | u3ble. | um4bly |
| 4cli | 4signa | $4{ }^{4} \mathrm{ply}$ | suit3 | te2ma2 | 3tine. | tricia | 13 ca | u1旡1 |
| -cors | -114e | s4pon | 84 l | tem3at | 3tini | 4trica | uci4b | ufiming |
| 4scopy | 4sily | spor4 | su2m | 3 tenan | 1 tio | 2trin | ucait | umor50 |
| scour5a | 2sin | 4spot | sum31 | 3 tene | ti5oc | tri4v | ucle3 | um2p |
| sica | s2ina | squal4l | su 2 n | 3 tend | tion5ee | ro5mi | 3 cr | unat4 |
| 485 d | 5sine. | 1 I | 8 c 2 r | 4 tenes | 5tiq | tron51 | u3cu | u2ne |
| 4 se . | 3ing | 2ss | 4 sv | 1 tent | ti3sa | 4trony | 14 cy | nfer |
| sela | 1sio | sisa | sw2 | ten4tag | 3tise | tro5phe | d5d | u1n1 |
| seast | 5sion | ssas3 | 48wo | 1 teo | tis4m | tro3sp | d3er | un4ia |
| sea5\% | sion5a | 8285c | 44y | te4p | ti5so | tro3v | d5est | unin |
| se2c3o | -12r | -3sel | 48yc | te5pe | tia4p | tru5i | udev4 | un5ish |
| 3sect | cir5a | 5seng | 3 syl | ter3c | Stistica | trus4 | u1dic | uni3v |
| 4a4ed | 181s | stses. | syn5o | 5 ter3d | ti3tl | 4t1:2 | ad3ied | un3e4 |
| se4d4* | 3aitio | s5set | sy5rin | 1 teri | ti4u | tisc | d3ies | andow |
| s5edl | 5siu | s1si | 1 ta | ter5ies | 1tiv | tsh4 | ud5is | ant3ab |
| se2g | 1siv | steie | 3 ta . | ter3is | tiv4a | teicm | u5dit | $n 4$ ter. |
| eeg3r | $58 i z$ | seifer | 2 tab | teri5za | 1tiz | $4 \mathrm{t} 3 \mathrm{t2}$ | 14don | unftes |
| 5ael | sk2 | sa5ily | ta5blea | 5 ternit | ti3za | t4tes | udiei | unu4 |
| selle | 4ske | 54s1 | 5taboliz | ter5v | ti3zon | t5to | u4du | un5y |
| 5 self | s3ket | s8411 | 4 taci | 4 tes. | 2 tl | ttu4 | u4ene | un5z |
| 5 selv | sk5ine | 84 sm | ta5do | 4tes: | t5la | 1 tu | uens4 | ufora |
| 4seme | sk5ing | sspend4 | 4 tal4 | t3ess. | tland | tula | uentto | u50e |
| se4mol | s 112 | ss2t . | taislo | teth5e | 3 tle . | tu3ar | uer4il | ulou |
| sen5at | s3lat | ssur5a | ta21 | 3teu | 3 tled | tu4bi | 3 ufa | u1pe |
| 4 senc | s21e | 855 | ta5la | 3 tex | 3 tles. | tud2 | u3f1 | uper5s |
| sen4d | slith5 | 2st. | talsen | 4tey | t5let. | 1tue | ugh3en | u5pia |
| J | $\checkmark$ | $\checkmark$ | $\sqrt{ }$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |  | $\checkmark$ |

feX82 HYPIIENATION PATTERNS


## Answers

```
moun-tain-ous vil-lain-ous
be-tray-al de-fray-al por-tray-al
hear-ken
ex-treme-ly su-preme-ly
tooth-aches
bach-e-lor ech-e-lon
riff-raff
anal-o-gous ho-mol-o-gous
gen-u-ine
any-place
co-a-lesce
fore-warn fore-word
de-spair
ant-arc-tic corn-starch
mast-odon
squirmed
```


## References

[1] Knuth, Donald E. TEX and METAFONT, New Directions in Typesetting. Digital Press, 1979.
[2] Webster's Third New International Dictionary. G. \& C. Merriam, 1981.
[3] Knuth, Donald E. The WEB System of Structured Documentation. Preprint, Stanford Computer Science Dept., September 1982.
[4] Knuth, Donald E. The Art of Computer Programming, Vol. 3, Sorting and Searching. Addison-Wesley, 1973.
[5] Standish, T. A. Data Structure Techniques. Addison-Wesley, 1980.
[6] Aho, A. V., Hopcroft, J. E., and Ullman, J. D. Algorithms and Data Structures. Addison-Wesley, 1982.
[7] Bloom, B. Space/time tradeoffs in hash coding with allowable errors. CACM 13, July 1970, 422-436.
[8] Carter, L., Floyd, R., Gill, J., Markowsky, G., and Wegman, M. Exact and approximate membership testers. Proc. 10th ACM SIGACT Symp., 1978, 5965.
[9] de la Briandais, Rene. File searching using variable length keys. Proc. Western Joint Computer Conf. 15, 1959, 295-298.
[10] Fredkin, Edward. Trie memory. CACM 3, Sept. 1960, 490-500.
[11] Trabb Pardo, Luis. Set representation and set intersection. Ph.D. thesis, Stanford Computer Science Dept., December 1978.
[12] Mehlhorn, Kurt. Dynamic binary search. SIAM J. Computing 8, May 1979, 175-198.
[13] Maly, Kurt. Compressed tries. CACM 19, July 1976, 409-415.
[14] Knuth, Donald E. TEX82. Preprint, Stanford Computer Science Dept., September 1982.
[15] Resnikoff, H. L. and Dolby, J. L. The nature of affixing in written English. Mechanical Translation 8, 1965, 84-89. Part II, June 1966, 23-33.
[16] The Merriam-Webster Pocket Dictionary. G. \& C. Merriam, 1974.
[17] Gorin, Ralph. SPELL.REG[UP,DOC] at SU-AI.
[18] Peterson, James L. Computer programs for detecting and correcting spelling errors. CACM 23, Dec. 1980, 673-687.
[19] Nix, Robert. Experience with a space-efficient way to store a dictionary. CACM 24, May 1981, 297-298.
[20] Morris, Robert and Cherry, Lorinda L. Computer detection of typographical errors. IEEE Trans. Prof. Comm. PC-18, March 1975, 54-64.
[21] Downey, P., Sethi, R., and Tarjan, R. Variations on the common subexpression problem. JACM 27, Oct. 1980, 758-771.
[22] Tarjan, R. E. and Yao, A. Storing a sparse table. CACM 22, Nov. 1979, 606-611.
[23] Zeigler, S. F. Smaller faster table driven parser. Unpublished manuscript, Madison Academic Computing Center, U. of Wisconsin, 1977.
[24] Aho, Alfred V. and Ullman, Jeffrey D. Principles of Compiler Design, sections 3.8 and 6.8. Addison-Wesley, 1977.
[25] Pfleeger, Charles P. State reduction in incompletely specified finite-state machines. IEEE Trans. Computers C-22, Dec. 1973, 1099-1102.
[26] Kohavi, Zvi. Switching and Finite Automata Theory, section 10-4. McGrawHill, 1970.
[27] Knuth, D. E., Morris, J. H., an i Pratt, V. R. Fast pattern matching in stringr. SIAM J. Computing 6, June 1977, 323-350.
[28] Aho, A. V. In R. V. Book (ed.), Formal Language Theory: Perspectives and Open Problems. Academic Press, 1980.
[29] Kucera, Henry and Francis, W. Nelson. Computational Analysis of Present-Day American English. Brown University Press, 1967.
[30] Research and Engineering Council of the Graphic Arts Industry. Proceedings of the 13th Annual Conference, 1963.
[31] Stevens, M. E. and Little, J. L. Automatic Typographic-Quality Typesetting Techniques: A State-of-the-Art Review. National Bureau of Standards, 1967.
[32] Berg, N. Edward. Electronic Composition, A Guide to the Revolution in Typesetting. Graphical Arts Technical Foundation, 1975.
[33] Rich, R. P. and Stone, A. G. Method for hyphenating at the end of a printed line. CACM 8, July 1965, 444-445.
[34] Wagner, M. R. The search for a simple hyphenation scheme. Bell Laboratories Technical Memorandum MM-71-1371-8.
[35] Gimpel, James F. Algorithms in Snobol 4. Wiley-Interscience, 1976.
[36] Ocker, Wolfgang A. A program to hyphenate English words. IEEE Trans. Prof. Comm. PC-18, June 1975, 78-84.
[37] Moitra, A., Mudur, S. P., and Narwekar, A. W. Design and analysis of a hyphenation procedure. Software Prac. Exper. 9, 1979, 325-337.
[38] Lindsay, R., Buchanan, B. G., Feigenbaum, E. A., and Lederberg, J. DENDRAL. McGraw-Hill, 1980.


[^0]:    - The work described in this thesis was inspired by the need for a word hyphenation routine as part of Don Knuth's $T_{\mathbb{E}} \mathrm{X}$ typesetting system [1]. This system was initially designed in order to typeset Prof. Knuth's seven-volume series of books, The Art of Computer Programming, when he became dissatisfied with the quality of computer typesetting done by his publisher. Since Prof. Knuth's books were to be a definitive treatise on computer science, he could not bear to see his scholarly work presented in an inferior manner, when the degradation was entirely due to the fact that the material had been typeset by a computer!

    Since then, $\mathrm{T}_{\mathrm{E}} \mathrm{X}$ (also known as Tau Epsilon Chi, a system for technical text) has gained wide popularity, and it is being adopted by the American Mathematical Society, the world's largest publisher of mathematical literature, for use in its journals. $\mathrm{T}_{\mathrm{E}} \mathrm{X}$ is distinctive among other systems for word processing/document preparation in its emphasis on the highest quality output, especially for technical material.

    One necessary component of the system is a computer-based algorithm for hyphenating English words. This is part of the paragraph justification routine, and it is intended to eliminate the need for the user to specify word division points explicitly when they are necessary for good paragraph layout. Hyphenation occurs relatively infrequently in most book-format printing, but it becomes rather critical in narrow-column formats such as newspaper printing. Insufficient attention paid to this aspect of layout results in large expanses of unsightly white space, or (even worse) in words split at inappropriate points, e.g. new-spaper.

    Hyphenation algorithms for existing typesetting systems are usually either rulebased or dietionary-based. Rule-based algorithms rely on a set of division rules such as given for English in the preface of Webster's Unabridged Dictionary [2]. These include recognition of common prefixes and suffixes, splitting between double consonants, and other more specialized rules. Some of the "rules" are not particularly

