Chapter 28

Huffman Coding

By Sariel Har-Peled, March 8, 2019[①]

28.1. Huffman coding

A ***binary code*** assigns a string of 0s and 1s to each character in the alphabet. A code assigns for each symbol in the input a codeword over some other alphabet. Such a coding is necessary, for example, for transmitting messages over a wire, were you can send only 0 or 1 on the wire (i.e., for example, consider the good old telegraph and Morse code). The receiver gets a binary stream of bits and needs to decode the message sent. A prefix code, is a code where one can decipher the message, a character by character, by reading a prefix of the input binary string, matching it to a code word (i.e., string), and continuing to decipher the rest of the stream. Such a code is a ***prefix code***.

A binary code (or a prefix code) is ***prefix-free*** if no code is a prefix of any other. ASCII and Unicode's UTF-8 are both prefix-free binary codes. Morse code is a binary code (and also a prefix code), but it is not prefix-free; for example, the code for S ($\cdots$) includes the code for E ($\cdot$) as a prefix. (Hopefully the receiver knows that when it gets $\cdots$ that it is extremely unlikely that this should be interpreted as EEE, but rather S.

Any prefix-free binary code can be visualized as a binary tree with the encoded characters stored at the leaves. The code word for any symbol is given by the path from the root to the corresponding leaf; 0 for left, 1 for right. The length of a codeword for a symbol is the depth of the corresponding leaf. Such trees are usually referred to as ***prefix trees*** or ***code trees***.
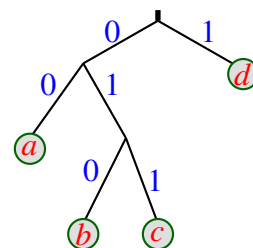
The beauty of prefix trees (and thus of prefix odes) is that decoding is easy. As a concrete example, consider the tree on the right. Given a string '010100', we can traverse down the tree from the root, going left if get a '0' and right if we get '1'. Whenever we get to a leaf, we output the character output in the leaf, and we jump back to the root for the next character we are about to read. For the example '010100', after reading '010' our traversal in the tree leads us to the leaf marked with 'b', we jump back to the root and read the next input digit, which is '1', and this leads us to the leaf marked with 'd', which we output, and jump back to the root. Finally, '00' leads us to the leaf marked by 'a', which the algorithm output. Thus, the binary string '010100' encodes the string "bda".

Suppose we want to encode messages in an *n*-character alphabet so that the encoded message is as short as possible. Specifically, given an array frequency counts $f[1\ldots n]$, we want to compute a prefix-free binary code that minimizes the total encoded length of the message. That is we would like

char	frequency		char	frequency		char	frequency		char	frequency
newline	16, 492		'0'	20		'A'	48,165		'N'	42,380
space	130,376		'1'	61		'B'	8,414		'O'	46,499
'!'	955		'2'	10		'C'	13,896		'P'	9,957
'"'	5,681		'3'	12		'D'	28,041		'Q'	667
'$'	2		'4'	10		'E'	74,809		'R'	37,187
'%'	1		'5'	14		'F'	13,559		'S'	37,575
'"'	1,174		'6'	11		'G'	12,530		'T'	54,024
'('	151		'7'	13		'H'	38,961		'U'	16,726
')'	151		'8'	13		'I'	41,005		'V'	5,199
'*'	70		'9'	14		'J'	710		'W'	14,113
','	13,276		':'	267		'K'	4,782		'X'	724
'_'	2,430		';'	1,108		'L'	22,030		'Y'	12,177
'.'	6,769		'?'	913		'M'	15,298		'Z'	215

' '	182
';'	93
'@'	2
'/'	26

Figure 28.1: Frequency of characters in the book "A tale of two cities" by Dickens. For the sake of brevity, small letters were counted together with capital letters.

char	frequency	code		char	frequency	code		char	frequency	code
'A'	48165	1110		'I'	41005	1011		'R'	37187	0101
'B'	8414	101000		'J'	710	1111011010		'S'	37575	1000
'C'	13896	00100		'K'	4782	11110111		'T'	54024	000
'D'	28041	0011		'L'	22030	10101		'U'	16726	01001
'E'	74809	011		'M'	15298	01000		'V'	5199	1111010
'F'	13559	111111		'N'	42380	1100		'W'	14113	00101
'G'	12530	111110		'O'	46499	1101		'X'	724	1111011011
'H'	38961	1001		'P'	9957	101001		'Y'	12177	111100
				'Q'	667	1111011001		'Z'	215	1111011000

Figure 28.2: The resulting prefix code for the frequencies of Figure 28.1. Here, for the sake of simplicity of exposition, the code was constructed only for the A—Z characters.

to compute a tree $\mathcal{T}$ that minimizes

$$\text{cost}(\mathcal{T}) = \sum_{i=1}^{n} f[i] * \text{len}\big(\text{code}(i)\big),$$

where code(i) is the binary string encoding the ith character and len(s) is the length (in bits) of the binary string s.

As a concrete example, consider Figure 28.1, which shows the frequency of characters in the book "A tale of two cities", which we would like to encode. Consider the characters 'E' and 'Q'. The first appears > 74,000 times in the text, and other appears only 667 times in the text. Clearly, it would be logical to give 'E', the most frequent letter in English, a very short prefix code, and a very long (as far as number of bits) code to 'Q'.

A nice property of this problem is that given two trees for some parts of the alphabet, we can easily put them together into a larger tree by just creating a new node and hanging the trees from this common node. For example, putting two characters together, we have the following.
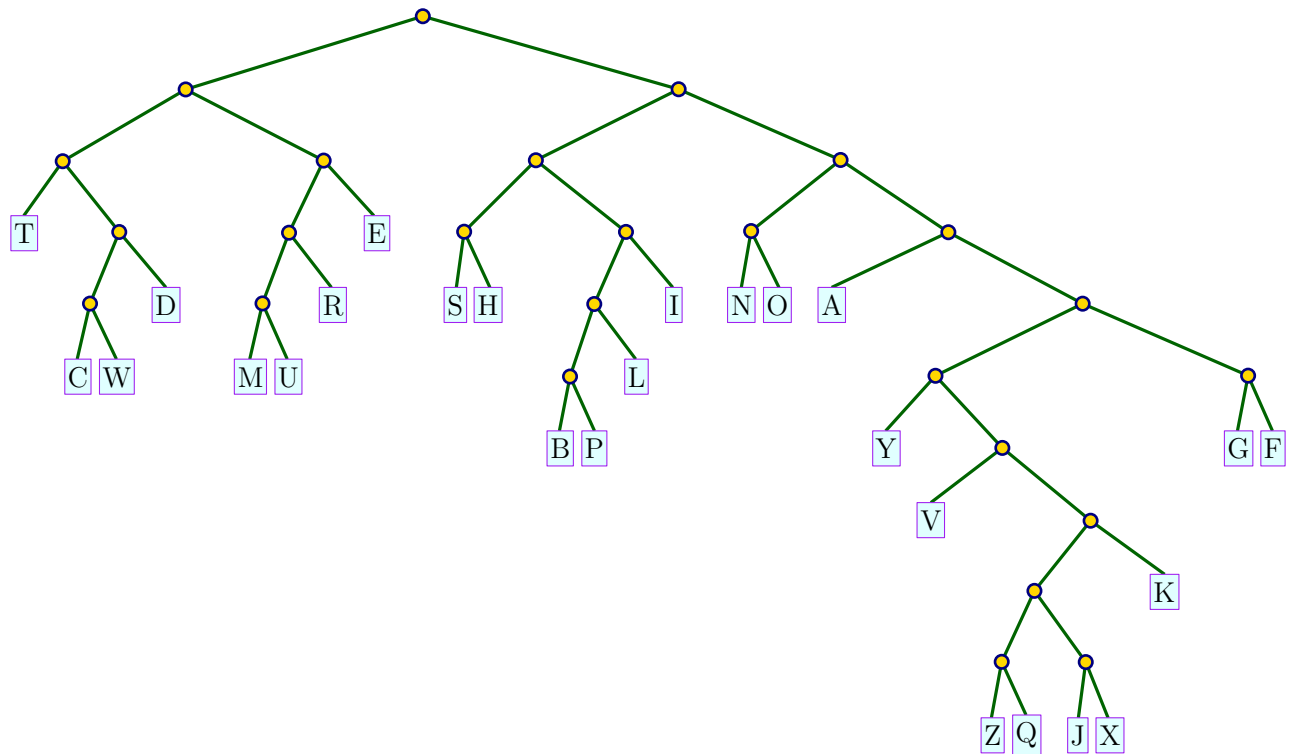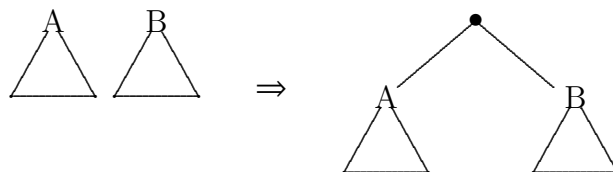
Figure 28.3: The Huffman tree generating the code of Figure 28.2.

Similarly, we can put together two subtrees.

28.1.1. The algorithm to build Hoffman's code

This suggests a simple algorithm that takes the two least frequent characters in the current frequency table, merge them into a tree, and put the merged tree back into the table (instead of the two old trees). The algorithm stops when there is a single tree. The intuition is that infrequent characters would participate in a large number of merges, and as such would be low in the tree – they would be assigned a long code word.

This algorithm is due to David Huffman, who developed it in 1952. Shockingly, this code is the best one can do. Namely, the resulting code is *asymptotically* gives the best possible compression of the data (of course, one can do better compression in practice using additional properties of the data and careful hacking). This ***Huffman coding*** is used widely and is the basic building block used by numerous other compression algorithms.

To see how such a resulting tree (and the associated code) looks like, see Figure 28.2 and Figure 28.3.

28.1.2. Analysis

Lemma 28.1.1. *Let $\mathcal{T}$ be an optimal code tree. Then $\mathcal{T}$ is a full binary tree (i.e., every node of $\mathcal{T}$ has either 0 or 2 children).*

In particular, if the height of $\mathcal{T}$ is d, then there are leafs nodes of height d that are sibling.

Proof: If there is an internal node in $\mathcal{T}$ that has one child, we can remove this node from $\mathcal{T}$, by connecting its only child directly with its parent. The resulting code tree is clearly a better compressor, in the sense of Eq. (28.1).

As for the second claim, consider a leaf u with maximum depth d in $\mathcal{T}$, and consider it parent $v = \overline{p}(u)$. The node v has two children, and they are both leafs (otherwise u would not be the deepest node in the tree), as claimed. ∎

Lemma 28.1.2. *Let x and y be the two least frequent characters (breaking ties between equally frequent characters arbitrarily). There is an optimal code tree in which x and y are siblings.*

Proof: More precisely, there is an optimal code in which x and y are siblings and have the largest depth of any leaf. Indeed, let $\mathcal{T}$ be an optimal code tree with depth d. The tree $\mathcal{T}$ has at least two leaves at depth d that are siblings, by Lemma 28.1.1.

Now, suppose those two leaves are not x and y, but some other characters α and β. Let $\mathcal{T}'$ be the code tree obtained by swapping x and α. The depth of x increases by some amount Δ, and the depth of α decreases by the same amount. Thus,

$$\text{cost}(\mathcal{T}') = \text{cost}(\mathcal{T}) - \left(f[\alpha] - f[x]\right)\Delta.$$

By assumption, x is one of the two least frequent characters, but α is not, which implies that $f[\alpha] > f[x]$. Thus, swapping x and α does not increase the total cost of the code. Since $\mathcal{T}$ was an optimal code tree, swapping x and α does not decrease the cost, either. Thus, $\mathcal{T}'$ is also an optimal code tree (and incidentally, $f[\alpha]$ actually equals $f[x]$). Similarly, swapping y and b must give yet another optimal code tree. In this final optimal code tree, x and y as maximum-depth siblings, as required. ∎

Theorem 28.1.3. *Huffman codes are optimal prefix-free binary codes.*

Proof: If the message has only one or two different characters, the theorem is trivial. Otherwise, let $f[1 \ldots n]$ be the original input frequencies, where without loss of generality, $f[1]$ and $f[2]$ are the two smallest. To keep things simple, let $f[n+1] = f[1] + f[2]$. By the previous lemma, we know that some optimal code for $f[1..n]$ has characters 1 and 2 as siblings. Let $\mathcal{T}_{\text{opt}}$ be this optimal tree, and consider the tree formed by it by removing 1 and 2 as it leaves. We remain with a tree $\mathcal{T}'_{\text{opt}}$ that has as leafs the characters $3, \ldots, n$ and a "special" character $n + 1$ (which is the parent of 1 and 2 in $\mathcal{T}_{\text{opt}}$) that has frequency $f[n+1]$. Now, since $f[n+1] = f[1] + f[2]$, we have

$$
\begin{aligned}
\text{cost}\left(\mathcal{T}_{\text{opt}}\right) &= \sum_{i=1}^{n} f[i]\text{depth}_{\mathcal{T}_{\text{opt}}}(i) \\
&= \sum_{i=3}^{n+1} f[i]\text{depth}_{\mathcal{T}_{\text{opt}}}(i) + f[1]\text{depth}_{\mathcal{T}_{\text{opt}}}(1) + f[2]\text{depth}_{\mathcal{T}_{\text{opt}}}(2) - f[n+1]\text{depth}_{\mathcal{T}_{\text{opt}}}(n+1) \\
&= \text{cost}\left(\mathcal{T}'_{\text{opt}}\right) + \left(f[1] + f[2]\right)\text{depth}\left(\mathcal{T}_{\text{opt}}\right) - \left(f[1] + f[2]\right)\left(\text{depth}\left(\mathcal{T}_{\text{opt}}\right) - 1\right) \\
&= \text{cost}\left(\mathcal{T}'_{\text{opt}}\right) + f[1] + f[2]. \quad\quad\quad\quad (28.1)
\end{aligned}
$$

4

This implies that minimizing the cost of $\mathcal{T}_{\text{opt}}$ is equivalent to minimizing the cost of $\mathcal{T}'_{\text{opt}}$. In particular, $\mathcal{T}'_{\text{opt}}$ must be an optimal coding tree for $f[3\ldots n+1]$. Now, consider the Huffman tree $\mathcal{T}'_H$ constructed for $f[3,\ldots,n+1]$ and the overall Huffman tree $\mathcal{T}_H$ constructed for $f[1,\ldots,n]$. By the way the construction algorithm works, we have that $\mathcal{T}'_H$ is formed by removing the leafs of 1 and 2 from $\mathcal{T}$. Now, by induction, we know that the Huffman tree generated for $f[3,\ldots,n+1]$ is optimal; namely, $\text{cost}\left(\mathcal{T}'_{\text{opt}}\right) = \text{cost}\left(\mathcal{T}'_H\right)$. As such, arguing as above, we have

$$\text{cost}(\mathcal{T}_H) = \text{cost}\left(\mathcal{T}'_H\right) + f[1] + f[2] = \text{cost}\left(\mathcal{T}'_{\text{opt}}\right) + f[1] + f[2] = \text{cost}(\mathcal{T}_{\text{opt}}),$$

by Eq. (28.1). Namely, the Huffman tree has the same cost as the optimal tree. ∎

28.1.3. What do we get

For the book "A tale of two cities" which is made out of 779,940 bytes, and using the above Huffman compression results in a compression to a file of size 439,688 bytes. A far cry from what `gzip` can do (301,295 bytes) or `bzip2` can do (220,156 bytes!), but still very impressive when you consider that the Huffman encoder can be easily written in a few hours of work.

(This numbers ignore the space required to store the code with the file. This is pretty small, and would not change the compression numbers stated above significantly.

28.1.4. A formula for the average size of a code word

Assume that our input is made out of n characters, where the ith character is p_i fraction of the input (one can think about p_i as the probability of seeing the ith character, if we were to pick a random character from the input).

Now, we can use these probabilities instead of frequencies to build a Huffman tree. The natural question is what is the length of the codewords assigned to characters as a function of their probabilities?

In general this question does not have a trivial answer, but there is a simple elegant answer, if all the probabilities are power of 2.

Lemma 28.1.4. *Let $1,\ldots,n$ be n symbols, such that the probability for the ith symbol is p_i, and furthermore, there is an integer $l_i \geq 0$, such that $p_i = 1/2^{l_i}$. Then, in the Huffman coding for this input, the code for i is of length l_i.*

Proof: The proof is by easy induction of the Huffman algorithm. Indeed, for $n = 2$ the claim trivially holds since there are only two characters with probability $1/2$. Otherwise, let i and j be the two characters with lowest probability. It must hold that $p_i = p_j$ (otherwise, $\sum_k p_k$ can not be equal to one). As such, Huffman's merges this two letters, into a single "character" that have probability $2p_i$, which would have encoding of length $l_i - 1$, by induction (on the remaining $n - 1$ symbols). Now, the resulting tree encodes i and j by code words of length $(l_i - 1) + 1 = l_i$, as claimed. ∎

In particular, we have that $l_i = \lg 1/p_i$. This implies that the average length of a code word is

$$\sum_i p_i \lg \frac{1}{p_i}.$$

If we consider X to be a random variable that takes a value i with probability p_i, then this formula is

$$\mathbb{H}(X) = \sum_i \mathbb{P}[X = i] \lg \frac{1}{\mathbb{P}[X = i]},$$

which is the ***entropy*** of X.

28.2. Bibliographical notes

This part of the class notes is based on Jeff Erickson class notes.

Bibliography