UNIVERSAL SOURCE CODING — LEMPEL–ZIV ALGORITHM

Our experience with data compression so far has been of the following type: We are given the statistical description of an information source, we then try to design a system which will represent the data produced by this source efficiently.

In this note we depart from this model, and consider a method which will represent a sequence efficiently without knowing by which means the sequence was produced. For this purpose, rather than assuming a statistical model for the sequence, it makes more sense to imagine that there is only a single sequence: an infinite string u which we wish to represent.

We will first consider the compressibility of an infinite string with a finite state machine. For our purposes, a finite state machine is a device that reads the input sequence one symbol at a time. Each symbol of the input sequence belongs to a finite alphabet $\mathcal{U}$ with J symbols ($J \geq 2$). The machine is in one of a finite number s of states before it reads a symbol, and goes to a new state determined by the old state and the symbol read. We will assume that the machine is in a fixed, known state z_1 before it reads the first input symbol. The machine also produces a finite string of binary digits (possibly the null string) after each input. This output string is again a function of the old state and the input symbol. That is, when the infinite sequence $u = u_1 u_2 \cdots$ is given as the input, the decoder produces $y = y_1 y_2 \cdots$, while visiting an infinite sequence of states $z = z_1 z_2 \cdots$, given by

$$y_k = f(z_k, u_k), \quad k \geq 1$$
$$z_{k+1} = g(z_k, u_k), \quad k \geq 1$$

where the function f takes values on the set of finite binary strings, so that each y_k is a (perhaps null) binary string. A finite segment $x_k x_{k+1} \cdots x_j$ of a sequence $x = x_1 x_2 \cdots$ will be denoted by x_k^j, and by an abuse of the notation, the functions f and g will be extended to indicate the output sequence and the final state. Thus, $f(z_k, u_k^j)$ will denote y_k^j and $g(z_k, u_k^j)$ will denote z_{j+1}.

To make the question of compressibility meaningful one has to require some sort of a 'unique decodability' condition on the finite state encoders. The decoder, given the description of the finite state machine that encoded the string, and the starting state z_1, but (of course) without the knowledge of the input string should be able to reconstruct the input string u from the output of the encoder y. A weaker requirement than this is the following: for any two distinct input sequences u_r^s and v_r^t, and for any z_r, either $f(z_r, u_r^s) \neq f(z_r, v_r^t)$ or $g(z_r, u_r^s) \neq g(z_r, v_r^t)$. An encoder satisfying this second requirement will be called *information lossless* (IL). It is clear that if an encoder is not IL, then there is no hope to recover the input from the output, and thus every 'uniquely decodable' encoder is IL. [1]

---

[1] However, as illustrated in Figure 1, an IL encoder is not necessarily uniquely decodable. Starting from state S, two distinct input sequences will leave the encoder in distinct states if they have different first symbols, otherwise they will lead to different output sequences. Thus, the above encoder is IL. Nevertheless, no decoder can distinguish between the input sequences `aaaa`$\cdots$ and `bbbb`$\cdots$ by observing the output $000 \cdots$.

A finite state machine with three states S, A and B. The notation $\mathtt{i}\,/\,output$ means that the machine produces *output* in response to the input $\mathtt{i}$. λ denotes the null output.

Figure 1: An IL encoder which is not uniquely decodable.

We will first derive a lower bound to the number of bits per input symbol *any* IL encoder will produce when encoding a string u. This lower bound will even apply to IL encoders which may have been designed with the advance knowledge about u. We will then show that a particular algorithm (the Lempel–Ziv algorithm) the design of which *does not* depend on u, does as well as this lower bound. That is to say, a machine that implements the LZ algorithm will compete well against any IL machine in compressing any u. (However, note that a machine that implements LZ will *not* be a finite state machine. The point is to show that LZ performs well — we are not interested in the fairness of the competition.)

We can now define the compressibility of an infinite string u. Given an IL encoder E, the compression ratio for the initial n symbols u_1^n of u with respect to this encoder is defined by

$$\rho_E(u_1^n) = \frac{1}{n}\,\mathrm{length}(y_1^n),$$

where $\mathrm{length}(y_1^n)$ is the length of the binary sequence y_1^n. (Note that since each y_i is a possibly null binary string $\mathrm{length}(y_1^n)$ may be more or less than n.) The minimum of $\rho_E(u_1^n)$ over the set of all IL encoders E with s or less states is denoted by $\rho_s(u_1^n)$. Observe that $\rho_s(u_1^n) \leq \lceil \log_2 J \rceil$. The compressibility of u with respect to the class of IL encoders with s or less states is then defined as

$$\rho_s(u) = \limsup_{n\to\infty} \rho_s(u_1^n).$$

Finally the compressibility of u with respect to IL encoders (or simply the *compressibility*) is defined as

$$\rho(u) = \lim_{s\to\infty} \rho_s(u).$$

Note that since $\rho_s(u)$ is non-increasing in s, the limit indeed exists.

Let us define $c(u_1^n)$ as the maximum number of distinct strings that u_1^n can be parsed into, including the null string. (Note that $c(u_1^n) \geq 1$.) It turns out that $c(u_1^n)$ plays a fundamental role in the compressibility of u.

Suppose u_1^n is parsed into $c \geq 1$ distinct strings. As any non-negative integer, c can be written as

$$c = \sum_{k=0}^{m-1} J^k + r$$

with $m \geq 0$ and $0 \leq r < J^m$. Since u_1^n is a concatenation of c distinct strings, the smallest value possible for n will result if these distinct strings as the shortest ones possible. Since there are J^k strings of length k, we see that

$$n \geq \sum_{k=0}^{m-1} k J^k + mr.$$

Since

$$\sum_{k=0}^{m-1} J^k = \frac{J^m - 1}{J - 1} \quad \text{and} \quad \sum_{k=0}^{m-1} k J^k = m \frac{J^m}{J - 1} - \frac{J}{J - 1} \frac{J^m - 1}{J - 1},$$

we see that

$$
\begin{aligned}
n &\geq m(c - r + 1/(J-1)) - (J/(J-1))(c - r) + mr \\
&\geq m(c + 1/(J-1)) - (J/(J-1))c \\
&\geq (m-2)c
\end{aligned}
$$

On the other hand, since $c < (J^{m+1} - 1)/(J - 1)$, we see that $J^{m+1} > (J-1)c + 1 > c$, which implies $m + 1 > \log_J(c)$ and see that

$$n > c(\log_J(c) - 3) = c \log_J(c/J^3),$$

and so

$$n > c(u_1^n)(\log_J(c(u_1^n)/J^3)). \tag{1}$$

Now we can state the following

THEOREM 1. *For any IL-encoder with s states,*

$$\text{length}(y_1^n) \geq c(u_1^n) \log_2(c(u_1^n)/(8s^2)). \tag{2}$$

Proof. Let u_1^n be parsed into $c = c(u_1^n)$ distinct words, $u = w_1 \ldots w_c$, and let c_{ij} be the number of words which find the machine in state i and leave it in state j. Because the machine is IL, the corresponding output sequences must be distinct, and their total length L_{ij} must satisfy (from (1), using $J = 2$ since y is a binary string)

$$L_{ij} \geq c_{ij} \log_2(c_{ij}/8).$$

The total length $\text{length}(y_1^n)$ is the sum of the L_{ij}'s, thus

$$\text{length}(y_1^n) \geq \sum_{1 \leq i,j \leq s} c_{ij} \log(c_{ij}/8).$$

Since $\sum_{i,j} c_{ij} = c(u_1^n)$, and since subject to this constraint the minimum of the right hand side occurs at $c_{ij} = c(u_1^n)/s^2$, (right hand side is a symmetric convex function) we get (2). $\qquad\square$

From (1) one can see that $c(u_1^n) = O(n/\log n)$. [Proof: Set $c' = c/J^3$ and $n' = n/J^3$. Note that (1) is equivalent to $c' \log_2 c' < n'$. Take n large enough so that $\sqrt{n'} \leq 2n'/\log_2 n'$. Now, either $c' < \sqrt{n'}$ or $c' \geq \sqrt{n'}$. In the first case $c' < 2n'/\log_2 n'$ by assumption. In the second, by (1), $c' < n'/\log_2 c' \leq n'/\log \sqrt{n'} = 2n'/\log_2 n'$. Thus, in either case $c' \leq 2n'/\log_2 n'$, and thus $c \leq 2n/\log_2(n/J^3)$.]

Using this and (2), we see that

$$\rho_s(u) \geq \limsup_{n\to\infty} \frac{1}{n} c(u_1^n) \log_2(c(u_1^n)/(8s^2))$$

$$= \limsup_{n\to\infty} \frac{1}{n} c(u_1^n) \log_2 c(u_1^n) - \lim_{n\to\infty} \frac{1}{n} c(u_1^n) \log_2(8s^2)$$

$$= \limsup_{n\to\infty} \frac{1}{n} c(u_1^n) \log_2 c(u_1^n)$$

and since the right hand side is independent of s,

$$\rho(u) \geq \limsup_{n\to\infty} \frac{1}{n} c(u_1^n) \log_2 c(u_1^n). \tag{3}$$

Now, let us describe the Lempel–Ziv algorithm. The algorithm proceeds by generating a dictionary for the source and constantly updating it. It starts up with a dictionary just consisting of the words of length 1, and operates in the following manner: When the dictionary has D words, each of its words is assigned a binary codeword of length $\lceil \log_2 D \rceil$ in lexicographic order. When a word in the dictionary is recognized in the input sequence, the encoder generates the binary codeword of that word on its output, and enlarges the dictionary by replacing the just recognized word with all its single letter extensions. The dictionary can be represented as a tree, whose leaves are the current dictionary entries. Figure 2 shows an example of the operation of the algorithm. Since the recognized words are encoded *before* the dictionary is modified, the decoder can keep track of the encoder's operation. Suppose that the algorithm parses the sequence u_1^n into $c_{lz}(u_1^n)$ words $w_1, \ldots, w_{c_{lz}}$. Then we can write:

$$u_1^n = \lambda w_1 w_2 \cdots w_{c_{lz}},$$

where λ denotes the null sequence. By construction, the first $c_{lz} - 1$ of the parses are distinct. (The last word $w_{c_{lz}}$ may not be distinct from the others.) If we concatenate the last two parses, and count in λ we get a parsing of u_1^n into $c_{lz}(u_1^n)$ distinct words. Thus $c_{lz}(u_1^n) \leq c(u_1^n)$. Since each parse extends the dictionary by $J - 1$ entries, the size of the dictionary at the end of parsing u_1^n is

$$1 + (J - 1)c_{lz}(u_1^n) \leq 1 + (J - 1)c(u_1^n) \leq Jc(u_1^n).$$

Thus, the number of bits the LZ algorithm emits after seeing u_1^n is

$$L_{lz}(y_1^n) \leq c_{lz}(u_1^n) \lceil \log_2(Jc(u_1^n)) \rceil \leq c_{lz}(u_1^n) \log_2(2Jc(u_1^n)) \leq c(u_1^n) \log_2(2Jc(u_1^n)).$$

Dividing by n, and taking the lim sup as n gets large we see

$$\limsup_{n\to\infty} \frac{1}{n} L_{lz}(y_1^n) \leq \limsup_{n\to\infty} \frac{1}{n} c(u_1^n)[\log_2 c(u_1^n) + \log_2(2J)]$$

$$= \limsup_{n\to\infty} \frac{1}{n} c(u_1^n) \log_2 c(u_1^n)$$

4

The parsing of the sequence `aaaccb` with the Lempel–Ziv algorithm. The figure shows the evolution of the dictionary. The sequence is parsed into the phrases `a`, `aa`, `c` and `cb`. Figure 2(a) shows the initial dictionary. In 2(b) we see the dictionary after reading `a`, 2(c) shows after `aaa` has been read, etc. At each stage one might assign each dictionary entry a fixed length binary codeword. If the assignment is done in lexicographic order, at stage (a) it will be $\{a \to 00, b \to 01, c \to 10\}$, at stage (b) $\{aa \to 000, ab \to 001, \ldots, c \to 100\}$, at stage (c) $\{aaa \to 000, aab \to 001, \ldots, cc \to 110\}$, and at stage (d) $\{aaa \to 0000, aab \to 0001, aac \to 0010, \ldots, cc \to 1000\}$, and the output sequence will be 00,000,110,0111. (Commas are put in to aid the reader, they will not appear at the output.)

Figure 2: Operation of the Lempel–Ziv algorithm

so that the LZ algorithm will achieve the lower bound previously derived (3) in the limit of $n \to \infty$. (However, the algorithm uses up infinite memory, since it keeps track of an ever growing tree.)

One can perhaps express the tradeoff we have seen as follows: suppose we want to compress an infinite string u, and we were given the choice of using the 'off the shelf" Lempel–Ziv, versus designing a machine tuned to u with a finite (but arbitrary) number of states. Then, we might as well pick the Lempel–Ziv: In the long run (i.e., for long strings) the Lempel–Ziv algorithm will do as well as the best finite state machine.

In particular, if one knew that the string u is the output of an information source which is stationary and ergodic, one could have designed a finite state machine that implements, for example, the Huffman algorithm designed for this source for a large enough block length that will compress the source output with high probability, arbitrary close to its entropy rate. Combined with the above paragraph we see that for such sources, the Lempel–Ziv algorithm will compress them to their entropy rate too.