Chapter 8

A Statistical Interpretation of Algorithmic Entropy

WE CAN REGARD a Turing machine as a kind of random generator, with a certain probability of outputting a given string. This suggests an alternative definition of the entropy of a string, more in line with the concepts of Shannon's statistical information theory. Fortunately, we are able to establish the equivalence of this new definition with our earlier one, at least up to an additive constant.

8.1 Statistical Algorithmic Entropy

Definition 8.1. *The statistical algorithmic entropy $h_T(s)$ of the string s with respect to the Turing machine T is defined by*

$$2^{-h_T(s)} = \sum_{p:T(p)=s} 2^{-|p|},$$

that is,

$$h_T(s) = -\lg \left(\sum_{p:T(p)=s} 2^{-|p|} \right).$$

We set

$$h(s) = h_U(s),$$

where U is our chosen universal machine.

Recall the convention that $\lg x$ denotes $\log_2 x$; all our logarithms will be taken to base 2.

Proposition 8.1. *For any Turing machine T, and any string s,*

$$h_T(s) \le H_T(s);$$

and in particular,

$$h(s) \le H(s).$$

Proof ▶. If T never outputs s then the sum is empty and so has value 0, giving

$$h_T(s) = \infty = H_T(s).$$

Otherwise, one of the programs that outputs s is the minimal program (for T) $p = \mu_T(s)$. Thus

$$2^{-h_T(s)} = \sum_{p:T(p)=s} 2^{-\|p\|} \ge 2^{-\|\mu_T(s)\|} = 2^{-H_T(s)},$$

and so

$$h_T(s) \le H_T(s).$$

Proposition 8.2.

$$h_T(s) \ge 0.$$

Proof ▶. Since

$$\{p : T(p) = s\} \subset \{p : T(p) \text{ defined}\},$$

it follows that

$$\sum_{p:T(p)=s} 2^{-|p|} \le \sum_{p:T(p) \text{ defined}} 2^{-|p|} \le 1,$$

by Kraft's Inequality 7.1.

Hence

$$h_T(s) = -\lg \sum_{p:T(p)=s} 2^{-|p|} \ge 0.$$

Proposition 8.3. *For any string $s \in \mathbb{S}$.*

$$h(s) < \infty.$$

Proof ▶. We can certainly construct a machine T which outputs a given string s without any input. Then

$$U(\langle T \rangle) = s.$$

Hence

$$H(s) \le |\langle T \rangle| < \infty,$$

and *a fortiori*

$$h(s) \le H(s) < \infty.$$

8.2 The Turing machine as random generator

Imagine the following scenario. Suppose we choose the input to T by successively tossing a coin. (We might employ a mad statistician for this purpose.) Thus if the current rule requires that T should read in an input bit, then a coin is tossed, and 1 or 0 is input according as the coin comes up heads or tails. The experiment ends if and when the machine halts.

The expectation that this results in a given string s being output is

$$\mathrm{P}_T(s) = \sum_{p:T(p)=s} 2^{-|s|}.$$

Recall that in Shannon's Information Theory, if an event e has probability p of occurring, then we regard the occurrence of e as conveying $-\lg p$ bits of information. (Thus occurrence of an unlikely event conveys more information than occurrence of an event that was more or less inevitable.)

In our case, the 'event' is the outputting of the string s. So the information conveyed by the string is just

$$h_T(s) = -\lg \mathrm{P}_T(s).$$

Summary

The statistical algorithmic entropy $h(s)$ gives us an alternative measure of the informational content of a string. Fortunately we shall be able to establish that it is equivalent to our earlier measure, up to the ubiquitous constant:

$$h(s) = H(s) + O(1).$$