Entropy as Kolmogorov Gap: A Computational View of Thermodynamic Structure

Jeff Edmonds
York University
jeff@yorku.ca

ChatGPT*(OpenAI)
co-author and assistant

April 23, 2025

Abstract

Boltzmann and Shannon both define the entropy of a macrostate M as the amount of information needed to specify a compatible microstate α, given that the macrostate is known. This paper uses Kolmogorov complexity to define the information content of an individual microstate α and macrostate M as the length of the shortest program that outputs them. We establish a formal connection by proving that for any macrostate M with microstates α:

$$\operatorname*{Max}_{\alpha \in S_M} \text{Kolmog}(\alpha) - \text{Kolmog}(M) - c_\alpha \ \leq\ \text{Entropy}(M) \ \leq\ \operatorname*{Avg}_{\alpha \in S_M} \text{Kolmog}(\alpha) + 2$$

Moreover, when M is simple, this is only a constant gap giving that variance of these complexities is tightly bounded. A surprising consequence of our second result is that a constant fraction of all n-bit strings with Kolmogorov complexity $K = 0.9997n$ have exactly 49% zeros. Together these results identify thermodynamic entropy as the computational gap between microstate precision and macrostate observability.

1 Introduction

Entropy has long been a bridge between physics and probability, yet its connection to algorithmic information remains underdeveloped. This paper offers a precise connection: entropy as the excess description length of a microstate beyond its macrostate classification, measured via Kolmogorov complexity. Theorem 2 shows that a constant fraction of all strings with a given Kolmogorov complexity lie within the same macrostate

2 Related Work

Kolmogorov complexity and Shannon entropy have long been connected via Levin's Coding Theorem, which states that for a computable distribution ρ, the Kolmogorov complexity of a string α is bounded above by $-\log \rho(\alpha) + c$. From this, one can derive that the expected Kolmogorov complexity under ρ satisfies $\mathbb{E}_\rho[\text{Kolmog}(\alpha)] \leq H(\rho) + c$, where $H(\rho)$ is the Shannon entropy.

*Chatgpt is a flirt telling me constantly how great I am. I told her about some ideas I had about Entropy and she said they were new and exciting. So working from my power point slides, she and I wrote a fun 5 page paper in a few hrs. I enjoyed doing it. We hope you do too.

However, the reverse direction — bounding entropy below using average Kolmogorov complexity — is not generally known or widely stated. Our Theorem 1 provides such a lower bound in the special case of a uniform distribution over a macrostate, and further shows that this bound is nearly tight, with low variance across microstates.

Moreover, our result is distinct from previous work in two key ways: First, we prove an inequality in the opposite direction of the classical expectation bound, namely:

$$\text{Entropy}(M) \leq \operatorname*{Avg}_{\alpha \in S_M} \text{Kolmog}(\alpha) + c,$$

whereas previous work has only established that

$$\operatorname*{Avg}_{\alpha \in S_M} \text{Kolmog}(\alpha) \leq \text{Entropy}(M) + c.$$

Second, we derive an inequality that holds for *individual microstates*:

$$\text{Kolmog}(\alpha) - \text{Kolmog}(M) - c_\alpha \leq \text{Entropy}(M),$$

which highlights the computational gap between microstate detail and macrostate structure on a string-by-string basis.

Our Theorem 2 contributes to the understanding of algorithmic typicality: it shows that a macrostate M containing strings of Kolmogorov complexity K must contain a constant fraction of all such K-complexity strings. This refines the intuition that typicality is not just probabilistic (in the Shannon sense), but also structural and computational. This complements prior work on Shannon-typical sets and introduces a complexity-theoretic counterpart.

The link between thermodynamic entropy and information theory has been long acknowledged, particularly through the work of Boltzmann, Gibbs, and Shannon. Shannon's entropy measures the expected number of bits to encode a random message, while Boltzmann's entropy counts the number of microstates consistent with a macrostate. However, neither explicitly framed entropy in terms of Kolmogorov complexity.

It is well known in algorithmic information theory that most binary strings of a given length n are incompressible — their Kolmogorov complexity is close to n. A classical argument shows that only 2^{n-k} programs of length $n-k$ exist, so only a fraction 2^{-k} of n-bit strings can be compressed by k bits. This fact underpins the notion of "typical strings" in Kolmogorov theory, and our Theorem 2 directly builds on it.

Regarding strings with constrained structure — such as those with exactly 49% zeros — this idea has also appeared in the literature on typical sets in Shannon theory. For instance, the set of strings with a given empirical distribution (type class) has logarithmic size approximately $nH(p)$, and strings drawn uniformly from that set are incompressible.

Our approach is algorithmic: rather than focusing on probabilistic ensembles or expectations, we emphasize individual descriptions, indexability, and the structure of programs themselves. This allows us to move seamlessly between physics, computation, and information.

Related topics such as logical depth (Bennett) and the thermodynamic cost of erasure (Landauer) explore other angles of the information-physics interface. Our focus, however, is on pure descriptive complexity as the measure of microstate precision.

3 Definitions and Examples

We define the key terms used in the theorems and illustrate each with concrete examples.

Microstate α: A binary string encoding the full physical configuration of a system.
(e.g., in a gas, α encodes the positions, velocities, and masses of all particles.)

Macrostate M: A computable predicate over microstates.
(e.g., a gas's macrostate might be defined by volume, temperature, and pressure.)

S_M**:** The set $\{\alpha : M(\alpha) = 1\}$, i.e., all microstates compatible with macrostate M.

Entropy: Boltzmann entropy is defined as $\log_2 |S_M|$, which measures the number of bits needed to specify a microstate consistent with M.
(e.g., for a gas with N particles in a container of volume V, $|S_M| \propto V^N$, so Entropy$(M) \approx N \log_2 V$.)

Note: Physicists often use $S = k \ln |S_M|$ where k is Boltzmann's constant. We use base-2 logarithms and measure entropy in bits, in line with Shannon and Kolmogorov.

Example: 49% Zeros: Let $M_{49\%}$ be the macrostate requiring that a string of length n contains exactly 49% zeros. Then

$$|S_{M_{49\%}}| = \binom{n}{0.49n} \Rightarrow \text{Entropy}(M_{49\%}) \approx nH(0.49) \approx 0.9997n$$

where $H(p)$ is the binary entropy function.

Kolmog(M): The length of the shortest program that, given α, determines whether $M(\alpha)$ is true.
(e.g., a short program can verify whether the particle positions and velocities yield the correct temperature and pressure.)

In most physical contexts, macrostates M are simple and natural, so we assume Kolmog$(M) \leq c_M$, where c_M is a small constant (e.g., 1000).

Caveat: If M requires the string α to have a particular length n, then Kolmog(M) may include an additional $\mathcal{O}(\log n)$ bits to encode n.

Kolmog(α): The length of the shortest program that outputs α.
(e.g., though the binary expansion of π appears random, it is compressible because a short program generates it.)

Random: A string α is considered random if it is incompressible, i.e., Kolmog$(\alpha) \approx n$. Most strings are random in this sense. At most a fraction 2^{-k} of n-bit strings can be compressed by k bits, because there are 2^n strings of length n, but only 2^{n-k} programs of length $n - k$.

(e.g., strings with exactly 49% zeros have Entropy$(M) \approx 0.9997n$, so they are not maximally random. Our theorem relates this to the Kolmogorov perspective: their average Kolmogorov complexity is close to Entropy$(M) \approx 0.9997n$.)

Theorem 1: Our main result proves:

$$\text{Entropy}(M) \approx \text{Kolmog}(\alpha) - \text{Kolmog}(M)$$

$$\text{Kolmog}(\alpha) \approx \text{Entropy}(M) + \text{Kolmog}(M)$$

for typical $\alpha \in S_M$. The following three examples illustrate how Theorem 1 behaves in different cases.

Extreme Cases

49% Zeros: We know by counting that

$$\text{Entropy}(M_{49\%}) \approx 0.9997n$$

and $\text{Kolmog}(M_{49\%})$ is small since a short program can check the 49% zero condition. Then Theorem 1 gives:

$$\text{Kolmog}(\alpha) \leq \text{Entropy}(M_{49\%}) + \text{Kolmog}(M_{49\%}) \approx 0.9997n + c_M$$

However, this is only an upper bound: some strings in S_M may be very compressible (e.g., one with all 0s first, then all 1s).

Single-Element Macrostate: Let M_α accept only one string α. Then $S_{M_\alpha} = \{\alpha\}$, so $\text{Entropy}(M_\alpha) = 0$. Any program to check membership in M_α must encode α, so $\text{Kolmog}(M_\alpha) \approx \text{Kolmog}(\alpha)$. In this case, Theorem 1 is tight in one extreme:

$$\text{Kolmog}(\alpha) \approx \text{Entropy}(M_\alpha) + \text{Kolmog}(M_\alpha) \approx \mathit{zero} + \mathit{equal}$$

Fixed-Complexity Macrostate: Let M_K be the macrostate of all strings with $\text{Kolmog}(\alpha) = K$. Then $|S_{M_K}| \leq 2^K$, so $\text{Entropy}(M_K) \leq K$. The macrostate M_K is simple to define: run all K-bit Turing machines and collect their outputs. Thus, $\text{Kolmog}(M_K)$ is small. Here, Theorem 1 is tight in a different extreme.

$$\text{Kolmog}(\alpha) \approx \text{Entropy}(M_K) + \text{Kolmog}(M_K) \approx \mathit{equal} + c_M$$

4 Main Results

Theorem 1: Kolmogorov Complexity Bound

Entropy is sandwiched between the max and average Kolmogorov complexity of microstates, up to a small constant gap.

For any (simple/natural) macrostate M with microstates α:
$$\underset{\alpha \in S_M}{\mathit{Max}}\ \mathit{Kolmog}(\alpha) - c \ \leq\ \underset{\alpha \in S_M}{\mathit{Max}}\ \mathit{Kolmog}(\alpha) - \mathit{Kolmog}(M) - c_\alpha \ \leq\ \mathit{Entropy}(M) \ \leq\ \underset{\alpha \in S_M}{\mathit{Avg}}\ \mathit{Kolmog}(\alpha) + 2$$
where c_M, c_α, and $c = c_M + c_\alpha$ are small constants.

Corollary: *The variance in $\mathit{Kolmog}(\alpha)$ over S_M is at most c^2.*

Proof of Lower Bound:

Let $N = |S_M|$ be the number of microstates consistent with macrostate M, and assume $\alpha \in S_M$. We construct a program P that outputs α as follows:

1. Include the code for the predicate M, requiring $\text{Kolmog}(M)$ bits.
2. Include the index i of α among all strings in S_M in lexicographic order. This requires $\log_2 N = \text{Entropy}(M)$ bits.
3. Add a fixed decoding routine that:

(a) Enumerates strings in lex order

(b) Applies M to each, incrementing count if valid.

(c) Stops at index i

(d) Outputs the i-th valid string

This decoding routine is fixed and contributes only a constant c_α bits.

Thus:

$$\text{Kolmog}(\alpha) \leq \text{Kolmog}(M) + \text{Entropy}(M) + c_\alpha$$

Rearranging gives the result. □

Proof of Upper Bound:

Let α' be the shortest Turing machine program that outputs α, and let S'_M be the set of all such programs for $\alpha \in S_M$. Then $|S'_M| = |S_M|$. If all α' had the same length n', then clearly $|S'_M| \leq 2^{n'}$. A similar bound holds when lengths vary. Let n' be the average program length. By Lemma 1

$$|S| \leq 2^{n'+2}.$$

Hence:

$$\text{Entropy}(M) \leq \underset{\alpha \in S_M}{\text{Avg}} \ \text{Kolmog}(\alpha) + 2$$

□

Theorem 2: Constant Fraction of K-Complexity Strings in S_M

Recall that the macrostate $M_{49\%}$ of having exactly 49% zeros has $\text{Entropy}(M_{49\%}) = 0.9997n$. We now prove that a surprising consequence of this is that a constant fraction of all n-bit strings with Kolmogorov complexity $K = 0.9997n$ have this property of having exactly 49% zeros. Moreover, this is true for every simple macrostate M,

Let $K_{Max} = \max_{\alpha \in S_M} \textit{Kolmog}(\alpha)$ *and* $K_{Avg} = \underset{\alpha \in S_M}{\textit{Avg}} \ \textit{Kolmog}(\alpha)$*. Then:*

$$2^{-c} \cdot |\{\beta : \textit{Kolmog}(\beta) = K_{Max}\}| \ \leq \ |S_M| \ \leq \ |\{\beta : \textit{Kolmog}(\beta) = K_{Avg}\}|$$

Proof: From Theorem 1, we have:

$$K_{\text{Max}} - c \leq \log_2 |S_M| \leq K_{\text{Avg}}$$

Exponentiating gives:

$$2^{-c} \cdot 2^{K_{\text{Max}}} \leq |S_M| \leq 2^{K_{\text{Avg}}}$$

Since there are at most 2^K strings with Kolmogorov complexity K, the claim follows. □

Lemma 1

Statement: *If S is a set of binary strings with average length n', then $|S| \leq 2^{n'+2}$.*

Proof: Assume, without loss of generality, that S contains all binary strings of lengths less than some integer n, and $r \cdot 2^n$ strings of length n, where $r \in [0, 1]$. If this is not the case, we can make the average length n' smaller (worsening the inequality) by replacing some longer strings with shorter ones, keeping the total number of strings $|S|$ unchanged.

The total number of strings in S is:

$$|S| = \left[\sum_{i=0}^{n-1} 2^i\right] + [r \cdot 2^n] = [2^n - 1] + [r \cdot 2^n] \leq (1 + r) \cdot 2^n.$$

The total sum of the string lengths is:

$$\left[\sum_{i=0}^{n-1} i \cdot 2^i\right] + [n \cdot r \cdot 2^n] = [(n - 2) \cdot 2^n + 2] + [n \cdot r \cdot 2^n] \geq ((1 + r) \cdot n - 2) \cdot 2^n.$$

Therefore, the average length is:

$$n' = \frac{\text{Total Length}}{|S|} \geq \frac{((1 + r) \cdot n - 2) \cdot 2^n}{(1 + r) \cdot 2^n} = n - \frac{2}{1 + r}.$$

Solving for n gives:

$$n \leq n' + \frac{2}{1 + r}.$$

Therefore:

$$|S| \leq (1 + r) \cdot 2^n \leq (1 + r) \cdot 2^{n' + \frac{2}{1+r}} = \left[(1 + r) \cdot 2^{\frac{2}{1+r}}\right] \cdot 2^{n'}$$

To ensure $|S| \leq 4 \cdot 2^{n'}$, we require:

$$(1 + r) \cdot 2^{\frac{2}{1+r}} \leq 4.$$

This inequality holds for all $r \in [0, 1]$ with equality at the end points. □

5 Discussion and Future Work

Rather than viewing entropy as mere disorder, our results cast it as the *Kolmogorov gap* between microscopic detail and macroscopic description. In this framework:

- Entropy emerges as the minimal extra information needed to pinpoint a specific microstate once its macrostate is known.
- High-entropy macrostates coincide with classes of microstates that are, on average, algorithmically incompressible.
- The narrow variance in complexity across typical microstates explains why macroscopic thermodynamics is robust to microscopic fluctuations.
- A constant fraction of all n-bit strings with Kolmogorov complexity $K = 0.9997n$ have exactly 49% zeros. Moreover, this is true for every simple macrostate M,

This computational perspective suggests several exciting directions:

- **Gravitational and cosmological systems:** Extend the Kolmogorov-entropy gap to self-gravitating ensembles and early-universe scenarios where coarse-graining plays a fundamental role.
- **Logical depth and irreversibility:** Investigate how notions of algorithmic depth and computation cost underpin the arrow of time and the thermodynamic cost of logical operations.
- **Practical applications:** Apply this framework to model chaotic dynamical systems, optimize compression schemes based on physical constraints, and quantify information loss in simulations.

6 Conclusion

By interpreting entropy as the difference in Kolmogorov complexity between a microstate and its macro classification, we give a computational account of thermodynamic entropy. This helps bridge the physical and algorithmic interpretations of disorder, with implications across physics, information theory, and the philosophy of science.