AN INTRODUCTION TO KOLMOGOROV COMPLEXITY

Stephen G. Simpson
Department of Mathematics
Pennsylvania State University
http://www.math.psu.edu/simpson/
simpson@math.psu.edu

NSF-DMS-0600823, NSF-DMS-0652637,
Grove Endowment, Templeton Foundation

Theory Seminar
Computer Science and Engineering
Pennsylvania State University
September 14, 2009

Abstract.

Let F be a finite mathematical object. For instance, F could be an integer, or a finite sequence of 0's and 1's, or a computer program, or a finite graph, or a finite group. Associated with F is a positive integer K(F), the Kolmogorov complexity of F, which measures the "amount of information" in F, or the minimum length of a "description" of F. More precisely, the Kolmogorov complexity of F is the minimum length of a computer program P which describes F in the sense that, if we run P with no inputs, then P eventually halts with output F. In this talk we define Kolmogorov complexity and prove some of its basic properties. We also survey some connections between Kolmogorov complexity and various mathematical topics. Among these topics are algorithmic randomness, Hausdorff dimension, diagonal noncomputability, degrees of unsolvability, tilings of the plane, and 2-dimensional symbolic dynamics.

A *bitstring* is a finite sequence of 0's and 1's. If ρ and σ are bitstrings, $\rho^\frown\sigma$ denotes the *concatenation*, ρ followed by σ. We write $|\sigma|$ = the length of σ. We use $\{0,1\}^*$ to denote the set of bitstrings.

A *machine* is a partial recursive function from bitstrings to bitstrings, $M :\subseteq \{0,1\}^* \to \{0,1\}^*$.

Note that the domain of M, dom(M), is a recursively enumerable set of bitstrings.

A *universal machine* is a machine U with the following property: For all machines M there exists a bitstring ρ such that $U(\rho^\frown\sigma) \simeq M(\sigma)$ for all bitstrings σ.

The existence of a universal machine is an easy consequence of the Enumeration Theorem for partial recursive functions.

The *Kolmogorov complexity* of a bitstring τ is defined as $\mathsf{C}(\tau) = \min\{|\sigma| \mid U(\sigma) \simeq \tau\}$ where U is a universal machine.

Note that $\mathsf{C}(\tau)$ is well defined up to $\pm O(1)$. In other words, if U_1 and U_2 are universal machines and $\mathsf{C}_i(\tau) = \min\{|\sigma| \mid U_i(\sigma) \simeq \tau\}$ for $i = 1, 2$, then $|\mathsf{C}_1(\tau) - \mathsf{C}_2(\tau)| \leq O(1)$ for all τ.

The *Kolmogorov complexity* of a positive integer n is defined as $\mathsf{C}(n) = \mathsf{C}(\langle \underbrace{1, \ldots, 1}_{n} \rangle)$.

Let F be a finite mathematical object.

For instance, F could be a rational number, a finite graph, a computer program, a formula, a finite group, a hereditarily finite set, etc.

The *Kolmogorov complexity* of F is defined as $\mathsf{C}(F) = \mathsf{C}(\#(F))$ where $\#(F)$ is the Gödel number of F.

Intuitively, $\mathsf{C}(F)$ is the size of the "shortest description" of F, measured in bits. In other words, the "amount of information" in F.

For instance, let τ be a bitstring of length 1,000,000,000.

If τ consists entirely of 1's, then τ is very easy to describe, so $C(\tau)$ is quite small.

However, the vast majority of bitstrings τ of length 1,000,000,000 are "random" and therefore hard to describe, in the sense that $C(\tau)$ is close to 1,000,000,000.

This is because, for a random τ of length 1,000,000,000, there is no good way to describe τ except by listing all of the bits of τ in order.

Some easily proved facts are:

1. $\mathsf{C}(|\tau|) \leq \mathsf{C}(\tau) + O(1)$.

2. $\mathsf{C}(\tau) \leq |\tau| + O(1)$.

3. $\mathsf{C}(\tau_1{}^\frown\tau_2) \leq 2\mathsf{C}(\tau_1) + 2\mathsf{C}(\tau_2) + O(1)$.

It would be nice to improve 3 to say that
$\mathsf{C}(\tau_1{}^\frown\tau_2) \leq \mathsf{C}(\tau_1) + \mathsf{C}(\tau_2) + O(1)$,
but unfortunately this is not the case.

In order to obtain such improvements, it is convenient to consider a variant concept.

A machine M is said to be *prefix-free* if $\mathrm{dom}(M)$ is prefix-free. In other words, there do not exist bitstrings σ_1 and σ_2 such that $M(\sigma_1)$ and $M(\sigma_2)$ are both defined and σ_1 is a proper initial segment of σ_2.

A *universal prefix-free machine* is
a prefix-free machine U with the following
property: For all prefix-free machines M
there exists a bitstring ρ such that
$U(\rho^\frown\sigma) \simeq M(\sigma)$ for all bitstrings σ.

The existence of a universal prefix-free
machine can be proved.

The *prefix-free Kolmogorov complexity* of
a bitstring τ is $\mathsf{K}(\tau) = \min\{|\sigma| \mid U(\sigma) \simeq \tau\}$
where U is a universal prefix-free machine.

Just as in the case of $\mathsf{C}(\tau)$, we can show that
$\mathsf{K}(\tau)$ is well defined up to $\pm O(1)$.

We can now improve 3 to 3′:

3′. $\mathsf{K}(\tau_1{}^\frown\tau_2) \leq \mathsf{K}(\tau_1) + \mathsf{K}(\tau_2) + O(1)$.

For positive integers n and finite
mathematical objects F, we define
$\mathsf{K}(n)$ and $\mathsf{K}(F)$ in the obvious way.

Comparing C ("plain complexity") with K ("prefix-free complexity"), we have:

$\mathsf{C}(\tau) \leq \mathsf{K}(\tau) + O(1)$.

$\mathsf{K}(\tau) \leq \mathsf{C}(\tau) + \mathsf{K}(\mathsf{C}(\tau)) + O(1)$.

$\mathsf{K}(\tau) \leq 2\mathsf{C}(\tau) + O(1)$.

$\mathsf{K}(\tau) \leq \mathsf{C}(\tau) + 2\log \mathsf{C}(\tau) + O(1)$.

$\mathsf{K}(\tau) \leq \mathsf{C}(\tau) + \log \mathsf{C}(\tau) + 2\log\log \mathsf{C}(\tau) + O(1)$.

Etc. Here log is the base 2 logarithm.

A useful technical lemma:

Let L be a recursively enumerable set of ordered pairs (m, τ) where m is a positive integer and τ is a bitstring.

Assume that $\sum_{(m,\tau)\in L} 1/2^m < \infty$.

Then $\mathsf{K}(\tau) \leq m + O(1)$ for all $(m, \tau) \in L$.

Complex infinite sequences.

Let X be an infinite sequence of 0's and 1's.

For each positive integer n, let $X \upharpoonright n$ be the bitstring consisting of the first n bits of X.

We say that X is *complex* if there exists a nontrivial, computable, lower bound on the Kolmogorov complexity of $X \upharpoonright n$.

More precisely, there exists a computable function $p : \mathbb{N} \to \mathbb{N}$ such that $\mathsf{K}(X \upharpoonright n) \geq p(n)$ for all n, and $\sup\{p(n) \mid n \in \mathbb{N}\} = \infty$.

Here $\mathbb{N}$ is the set of positive integers.

This means that, viewing X as a stream of binary information, the bits of X are at least somewhat difficult to predict.

In the above definition, it does not matter whether we use prefix-free complexity, K, or plain complexity, C.

Lemma. X is complex if and only if there exists a computable function $q : \mathbb{N} \to \mathbb{N}$ such that $\mathsf{K}(X \upharpoonright q(n)) \geq n$ for all n.

Proof. The idea is that $q = p^{-1}$. More precisely, $q(n) =$ the least m such that $p(m) \geq n$.

Again, we can use either K or C here.

An easy observation is:

Theorem. If X is complex, then X is not computable.

Proof. Assume that X is computable. Since q is computable, we clearly have $\mathsf{K}(X \upharpoonright q(n)) \leq \mathsf{K}(n) + O(1)$. It is also clear that $\mathsf{K}(n) \leq \log n + O(1)$. Combining these two observations, we have $\mathsf{K}(X \upharpoonright q(n)) \leq \log n + O(1)$, so let c be a constant such that $\mathsf{K}(X \upharpoonright q(n)) \leq \log n + c$ for all n. Let n be so large that $\log n + c < n$. Then $\mathsf{K}(X \upharpoonright q(n))$ is both $< n$ and $\geq n$, a contradiction, Q.E.D.

Diagonal nonrecursiveness.

Consider a (non-computable) function $f : \mathbb{N} \to \mathbb{N}$. We say that f is *diagonally nonrecursive* (DNR) if there is no computer program P which eventually halts with output $f(\#(P))$. We say that f is *recursively bounded* (RB) if there exists a computable function $q : \mathbb{N} \to \mathbb{N}$ such that $f(n) \leq q(n)$ for all n.

Theorem (Kjos-Hanssen/Merkle/Stephan, 2005)**.** The following problems have the same degree of unsolvability.

1. Find an $f : \mathbb{N} \to \mathbb{N}$ which is DNR and RB.
2. Find an $X \in \{0,1\}^\infty$ which is complex.

We omit the proof.

Note: Two problems S_1 and S_2 are said to have the same *degree of unsolvability* if any solution of S_1 can be used as a Turing oracle to compute a solution of S_2 and vice versa.

Hausdorff dimension and fractals.

Let X be an infinite sequence of 0's and 1's.

The *effective dimension of* X is defined as

$$\dim(X) = \liminf_{n\to\infty} \frac{\mathsf{K}(X \upharpoonright n)}{n}.$$

Note that $\dim(X)$ is a real number in the interval $\{r \mid 0 \leq r \leq 1\}$.

Again, we can use either K or C here.

Clearly $\dim(X) > 0$ if and only if there is a linear lower bound on $\mathsf{K}(X \upharpoonright n)$. Thus $\dim(X) > 0$ implies that X is complex. In fact, the complexity of X is rather high.

The space of all infinite sequences of 0's and 1's is denoted $\{0,1\}^\infty$. Topologically, $\{0,1\}^\infty$ is the product of infinitely many copies of the two-point space $\{0,1\}$.

Recall that a set $S \subseteq \{0,1\}^\infty$ is said to be *closed* if the limit of any convergent sequence of points in S belongs to S. Equivalently, S is the complement of an open set. Recall that an open set is just the union of a sequence of basic open sets in $\{0,1\}^\infty$.

A set $S \subseteq \{0,1\}^\infty$ is said to be *effectively closed* if it is the complement of a set which is *effectively open*, i.e., the union of a computable sequence of basic open sets.

Theorem (Lutz/Mayordomo/. . . , 2000). Assume that S is effectively closed. Then, the Hausdorff dimension of S is equal to $\sup\{\dim(X) \mid X \in S\}$.

A similar result can be proved for effectively closed sets in Euclidean space. This includes the familiar fractals in Euclidean space, e.g., the Sierpinski triangle. Thus we have:

Corollary. The Hausdorff dimension of a fractal is equal to the supremum of the effective dimensions of its points.

Algorithmic randomness.

Let X be an infinite sequence of 0's and 1's. We say that X is *random* if it behaves like a sequence of coin tosses. More precisely:

Definition (Martin-Löf, 1966)**.** Let μ be the fair coin probability measure on $\{0,1\}^\infty$. A point $X \in \{0,1\}^\infty$ is said to be *random* if $X \in \bigcup_{n=1}^\infty S_n$ whenever S_n, $n = 1, 2, 3, \ldots$ is a computable sequence of effectively closed sets such that $\mu(S_n) \geq 1 - 1/2^n$ for all n.

It can be shown that any X which is random in the sense of Martin-Löf passes all effective statistical tests for randomness. For instance, X obeys the *Strong Law of Large Numbers*

$$\lim_{n\to\infty} \frac{X_1 + \cdots + X_n}{n} = \frac{1}{2}$$

and the *Law of the Iterated Logarithm*

$$\limsup_{n\to\infty} \frac{|X_1 + \cdots + X_n - n/2|}{\sqrt{n \log\log n}} = \frac{1}{\sqrt{2}}$$

where log denotes the base e logarithm.

Theorem (Schnorr, 1975)**.**
The following properties are equivalent.

1. X is random in the sense of Martin-Löf.
2. $\mathrm{K}(X \upharpoonright n) \geq n - O(1)$ for all n.

Roughly speaking, Schnorr's Theorem means that X is random if and only if the bitstrings $X \upharpoonright n$ are "asymptotically incompressible".

In Schnorr's Theorem it is important that we are using prefix-free complexity, K, rather than plain complexity, C.

We omit the proof of Schnorr's Theorem.

Corollary. If X is random, then $\dim(X) = 1$.

However, easy examples show that the converse of the corollary is false.

Relativized notions.

Let A be a Turing oracle. Relative to A one may consider Kolmogorov complexity, effectively closed sets, and randomness.

Define $A \leq_{\mathsf{LR}} B$ to mean that every $X \in \{0,1\}^\infty$ which is B-random is A-random. In other words, B is at least as powerful as A in terms of detecting nonrandomness.

Define $A \leq_{\mathsf{LK}} B$ to mean that $\mathsf{K}^B(\tau) \leq \mathsf{K}^A(\tau) + O(1)$ for all bitstrings τ. In other words, B is at least as powerful as A in terms of compressing information.

Theorem (Kjos-Hanssen/Miller/Solomon, 2005)**.** $A \leq_{\mathsf{LR}} B$ if and only if $A \leq_{\mathsf{LK}} B$, if and only if every A-effectively closed set of positive measure includes a B-effectively closed set of positive measure.

Here again, it is important that we are using K rather than C.

Degrees of unsolvability.

Following Simpson 1999, let $\mathcal{E}_{\mathsf{w}}$ be the lattice of degrees of unsolvability associated with nonempty, effectively closed sets in $\{0,1\}^\infty$.

Many interesting degrees in $\mathcal{E}_{\mathsf{w}}$ are related to Kolmogorov complexity. For instance:

$\mathbf{d} = \deg(\{f \mid f \text{ is DNR}\})$.

$\mathbf{d}_C = \deg(\{f \mid f \text{ is DNR and } C\text{-bounded}\})$.

$\mathbf{d}_{\mathsf{REC}} = \deg(\{X \mid X \text{ is complex}\})$.

$\mathbf{q}_s = \deg(\{X \mid \dim(X) > s\})$.

$\mathbf{b}_1 = \deg(\{X \mid 0' \leq_{\mathsf{LR}} X\})$ where $0'$ is the halting problem for Turing machines.

$\mathbf{b}_\alpha = \deg(\{X \mid 0^{(\alpha)} \leq_{\mathsf{LR}} X\})$ where $0^{(\alpha)}$ is the αth iterate of the Turing jump operator.

$\mathbf{r}_1 = \deg(\{X \mid X \text{ is random}\})$.

$\mathbf{r}_2 = \deg(\{X \mid X \text{ is random relative to } 0'\})$.

$\mathbf{1} = \deg(\{f \mid f \text{ is DNR and 2-bounded}\})$
$= \deg(\{T \mid T \text{ is a completion of PA}\})$.

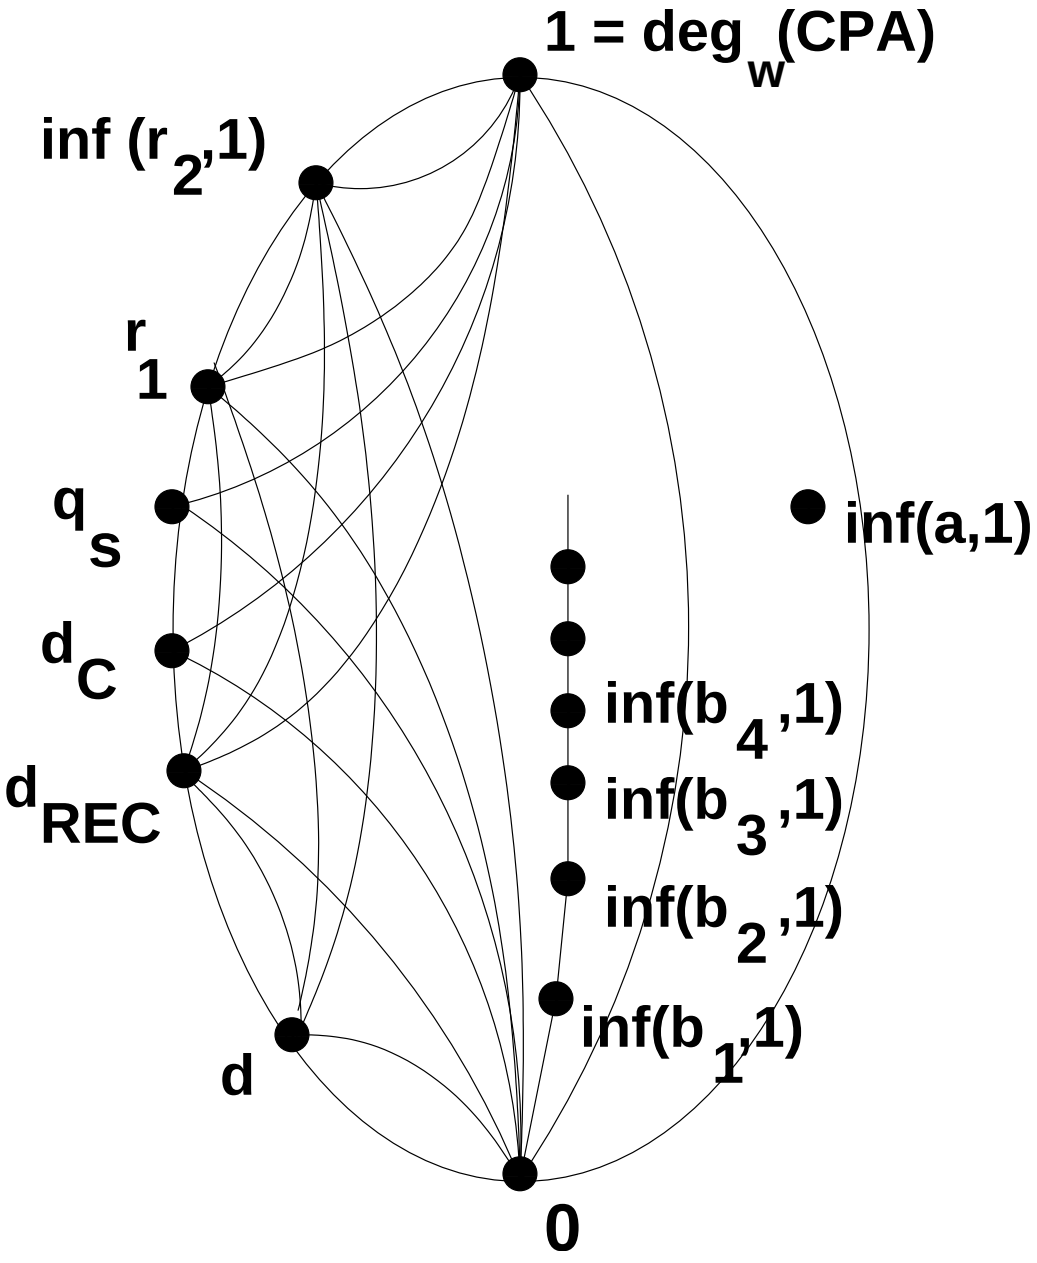

A picture of $\mathcal{E}_{\mathrm{w}}$. Here $\mathbf{a} =$ any r.e. degree, $\mathbf{r} =$ randomness, $\mathbf{b} =$ LR-reducibility, $\mathbf{q} =$ dimension, $\mathbf{d} =$ diagonal nonrecursiveness.

Tiling problems (Wang, 1961)**.**

Let F be a finite set of *tiles*, i.e., 1×1 squares with colored edges. Let P_F be the *tiling problem* associated with F, i.e, the problem of covering the Euclidean plane with disjoint copies of tiles from F in such a way that adjacent edges have matching colors.

Example. Let F be this set of four tiles:

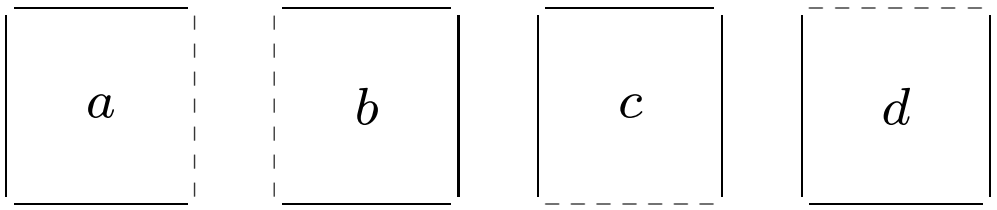

Then P_F is the problem of covering the plane with 2×1 and 1×2 rectangles

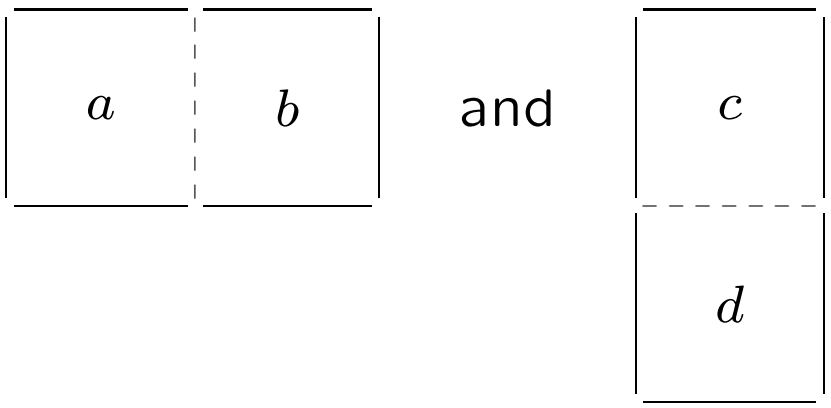

Example (continued).

The tiling problem P_F has many solutions:

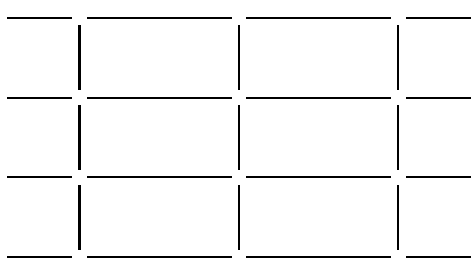

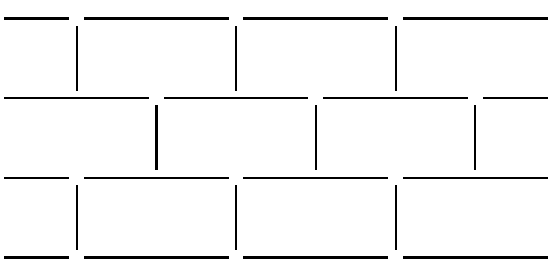

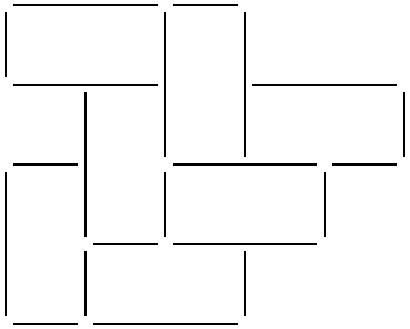

Here are some theorems showing that tiling problems can be very difficult to solve.

Theorem (Berger, 1966)**.** We can construct a tiling problem P_F which has solutions but no periodic solution.

Theorem (Myers, 1974)**.** We can construct a tiling problem P_F which has solutions but no computable solution.

Theorem (Durand/Romashchenko/Shen, 2008)**.** We can construct a tiling problem with the following property. P_F has solutions, but $\mathsf{K}(S) \geq O(n) - O(1)$ for any $n \times n$ square S in any solution of P_F.

Note: $O(n) - O(1)$ is best possible.

Here is another theorem in this vein.

Theorem (Simpson, 2007)**.**
Let P_F be a tiling problem which has solutions. Then, the degree of unsolvability of P_F belongs to $\mathcal{E}_{\mathrm{w}}$. Conversely, each degree in $\mathcal{E}_{\mathrm{w}}$ is the degree of a tiling problem.

My paper proving this result has been accepted for publication in the journal *Ergodic Theory and Dynamical Systems*.

Remark. The study of tiling problems is essentially the same as 2-dimensional symbolic dynamics. Given a tiling problem P_F, the solution set S_F is either empty or a 2-dimensional shift space of finite type. Conversely, each 2-dimensional shift space of finite type is equivalent to the solution set of a tiling problem.

Current research.

One of my research projects is
to study the relationship between
the degree of unsolvability of P_F and
the classical dynamical properties of
the dynamical system S_F.

An classically important invariant of
dynamical systems is entropy.

A less well-studied invariant is
degree of unsolvability.

Both of these invariants measure the
complexity of orbits in a dynamical system.
The entropy is an upper bound, while the
degree of unsolvability is a lower bound.

THE END.

THANK YOU!