

Certified Coverage of Exponential Search Spaces

J. R. Landers

September 2026

Idea. Consider a search problem whose naive state space has size 2^n . Rather than treating every state individually, suppose part of the space can be represented structurally, another part can be proposed for elimination by a learned model and then certified exactly, and only the residual must be searched explicitly. The central quantity is not the fraction removed, but the *growth of what remains uncovered*.

A coverage model

Let $\mathcal{S}_n$ be a nominal search space with

$$|\mathcal{S}_n| = 2^n.$$

We split the reduction of this space into two mechanisms.

First, a structural procedure identifies a set $K_n \subseteq \mathcal{S}_n$ of size

$$k(n) = |K_n|$$

that admits a compact representation. In the simplest logarithmic model, the $k(n)$ explicit states are replaced by a description of cost

$$C_{\text{comp}}(n) = \log_2 k(n).$$

This is deliberately an abstract assumption: it describes the cost once such a representation is available, not a universal property of arbitrary subsets.

Second, a learned classifier proposes regions of the remaining space for removal. Let $\delta(n)$ denote the total number of states covered by these proposals. The classifier is not trusted to delete states directly. Instead, an exact verifier certifies proposed regions. If $a(n) \in [0, 1]$ is the fraction of proposed mass that is successfully certified, then the safely eliminated mass is

$$\delta_c(n) = a(n)\delta(n).$$

Prediction error therefore affects efficiency, not correctness: uncertified proposals simply remain in the residual problem.

The explicit residual is

$$r(n) = 2^n - k(n) - a(n)\delta(n). \quad (1)$$

Let $P(n)$ collect the cost of discovering the structural representation, running the learned proposal mechanism, certifying its proposed regions, and managing the resulting representation. Under the simple assumption that explicit residual search costs $O(r(n))$, the total modeled cost is

$$T(n) = P(n) + r(n) + \log_2 k(n). \quad (2)$$

Thus

$$T(n) = P(n) + 2^n - k(n) - a(n)\delta(n) + \log_2 k(n). \quad (3)$$

The separation is important. A large $k(n)$ or $\delta(n)$ is useful only if the machinery that discovers, represents, and certifies those regions is itself cheap.

The certified coverage theorem

Define the total *certified coverage*

$$C(n) = k(n) + a(n)\delta(n),$$

so that $r(n) = 2^n - C(n)$.

Theorem 1 (Certified coverage criterion). *Assume $P(n) = \text{poly}(n)$, $1 \leq k(n) \leq 2^n$, and the residual is searched in time $O(r(n))$. Then*

$$T(n) = \text{poly}(n) \iff 2^n - C(n) = \text{poly}(n).$$

Equivalently, polynomial modeled runtime occurs exactly when

$$C(n) = 2^n - \text{poly}(n).$$

Proof. From (2),

$$T(n) = P(n) + (2^n - C(n)) + \log_2 k(n).$$

Because $k(n) \leq 2^n$,

$$0 \leq \log_2 k(n) \leq n.$$

The first and third terms are therefore polynomially bounded. Hence $T(n)$ is polynomial exactly when the residual $2^n - C(n)$ is polynomially bounded. $\square$

The theorem gives a sharp asymptotic warning: eliminating a very large *percentage* of an exponential space need not change its asymptotic order.

Corollary 1 (Fixed uncovered fraction). *If there exists a constant $\varepsilon > 0$ such that, for infinitely many n ,*

$$C(n) \leq (1 - \varepsilon)2^n,$$

then

$$r(n) \geq \varepsilon 2^n$$

and therefore the modeled runtime satisfies

$$T(n) = \Omega(2^n)$$

up to the stated residual-search assumption.

Thus 90%, 99%, or even 99.999% coverage is asymptotically insufficient if the uncovered fraction stays bounded away from zero. Polynomial behavior requires

$$\frac{C(n)}{2^n} = 1 - \frac{\text{poly}(n)}{2^n},$$

so the uncovered fraction must vanish exponentially fast, up to polynomial factors.

Differential form

The discrete complexity statement has a useful continuous interpolation. Treat n as a continuous scale parameter and differentiate (1):

$$\boxed{\frac{dr}{dn} = (\ln 2)2^n - k'(n) - a(n)\delta'(n) - a'(n)\delta(n).} \quad (4)$$

The four terms have a direct interpretation:

$$\underbrace{(\ln 2)2^n}_{\text{creation of raw states}} - \underbrace{k'}_{\text{structural coverage}} - \underbrace{a\delta'}_{\text{new certified proposals}} - \underbrace{a'\delta}_{\text{improving certification rate}}.$$

Differentiating the full cost gives

$$\boxed{\frac{dT}{dn} = (\ln 2)2^n - \left(1 - \frac{1}{k \ln 2}\right) k' - a\delta' - a'\delta + P'(n).} \quad (5)$$

For large k , the $1/(k \ln 2)$ correction is negligible, giving the approximate balance law

$$\boxed{T'(n) \approx (\ln 2)2^n - k' - a\delta' - a'\delta + P'(n).} \quad (6)$$

Equivalently, since $r(n) = 2^n - C(n)$,

$$r'(n) = (\ln 2)2^n - C'(n).$$

Integrating from a reference scale n_0 ,

$$\boxed{r(n) = r(n_0) + \int_{n_0}^n [(\ln 2)2^s - C'(s)] ds.} \quad (7)$$

This form makes the central picture especially transparent: polynomial behavior requires the cumulative gap between exponential state creation and certified coverage growth to remain only polynomial in n .

Why certification matters

A naive learned-pruning model might discard $\delta(n)$ states with error rate $e(n)$. If every erroneous deletion has to be recovered to preserve exactness, an error residue of order $e(n)\delta(n)$ appears. When $\delta(n) = \Theta(2^n)$, any fixed $e > 0$ leaves an exponential residue. To keep that residue polynomial would require

$$e(n) 2^n = \text{poly}(n), \quad \text{hence} \quad \boxed{e(n) \leq \frac{\text{poly}(n)}{2^n}.$$

Thus ordinary constant-error prediction is not enough for exact bulk deletion of an exponential region.

Certification avoids that failure mode. The learned model becomes a *proposal mechanism*: it suggests large regions that may be irrelevant, while an exact verifier alone has authority to delete them. A poor classifier can waste computation, but it cannot make the final algorithm incorrect. In this formulation, learning changes δ and the acceptance function a ; correctness is carried by verification.

The resulting abstract algorithm is

$$\begin{aligned} &\text{structural compression} \longrightarrow \text{learned proposals} \longrightarrow \text{exact certification} \\ &\hspace{10em} \longrightarrow \text{exact residual search} \end{aligned}$$

Scope and the remaining difficulty

The coverage criterion is a theorem about the model; it does not assert that every concrete search problem admits such coverage. The hard question is whether, for a given problem, one can simultaneously achieve

$$P(n) = \text{poly}(n) \quad \text{and} \quad 2^n - C(n) = \text{poly}(n).$$

Cardinality alone cannot answer this: exponentially many states can sometimes be described by a short rule. The substantive lower- or upper-bound problem is therefore about the *cost of discovering and certifying structure*, not merely the number of states represented by that structure.

General representation cost

The logarithmic term is the motivating special case, not an essential part of the coverage theorem. If a compressed region of size $k(n)$ has representation cost $\rho(n, k)$, write

$$T(n) = P(n) + r(n) + \rho(n, k(n)).$$

Whenever both $P(n)$ and $\rho(n, k(n))$ are polynomially bounded, the same proof gives

$$\boxed{T(n) = \text{poly}(n) \iff r(n) = 2^n - C(n) = \text{poly}(n).}$$

Thus the essential object is the uncovered residual; $\log_2 k$ is simply the strongest compression regime considered here. This also separates two questions that should not be conflated: how many states a representation covers, and how costly that representation is to discover and manipulate.

Summary. In the model developed here, exponential search is converted into a competition between state-space growth and certified coverage. Compression accounts for states represented collectively; learning proposes additional regions; certification makes those eliminations exact; and the residual determines the remaining explicit work. The governing condition is simple:

$$\boxed{\text{polynomial overhead} + \text{polynomial residual} \implies \text{polynomial modeled runtime}.}$$

The differential form records the same statement dynamically: coverage must track exponential state creation closely enough that their accumulated gap is only polynomial.