Maximal LZ78 Complexity Among Sturmian Sequences

Noble Number Extremality and the Epoch-Packing Theorem

Jetxel Alfredo
Independent Researcher
Scalar Resonance Research Program — March 2026

Abstract. *We investigate the LZ78 phrase count of Sturmian sequences — the symbolic dynamics of irrational circle rotations — as a function of the rotation number α. We introduce an epoch decomposition framework that partitions the LZ78 parsing into segments corresponding to the convergent denominators of the continued fraction (CF) expansion of α, and establish four principal results. First, we prove the Sturmian Exponent Theorem: for any Sturmian sequence, the LZ78 phrase count satisfies $c(\alpha, L) = \Theta(L^{2/3})$. The upper bound follows from a constrained optimization on the LZ78 trie using only the factor complexity $p(n) = n + 1$. The lower bound uses a bootstrap argument in which the upper bound constrains the trie depth, which feeds into the epoch recurrence to produce a matching lower bound. The exponent 2/3 is specific to LZ78 and differs from the $\sqrt{L}$ scaling of LZ77 factorization complexity for Sturmian words. Second, we prove that all noble numbers — irrationals whose CF expansion is eventually [1, 1, 1, ...] — achieve asymptotically identical LZ78 phrase counts, differing by at most a bounded additive constant (Noble Universality Theorem). Third, we prove that among Sturmian sequences with periodic CF [0; M, M, ...], the phrase count at any given length L is eventually strictly decreasing in M, with the Fibonacci word ($M = 1$, slope $1/\varphi$) achieving the eventual maximum (Epoch-Packing Theorem). Fourth, we prove that the Fibonacci word eventually maximizes the LZ78 phrase count among all Sturmian sequences (Eventual Maximality Theorem). The mechanism is an epoch-packing argument: the per-symbol phrase creation rate at any given convergent scale is monotonically decreasing in M, so the $M = 1$ sequence (Fibonacci word) generates the most phrases both by having the most convergent scales and by generating phrases at the highest rate per scale. The golden ratio $1/\varphi = [0; 1, 1, 1, \ldots]$, whose convergent denominators grow as φ^k (the slowest possible rate for any irrational), packs the most scales into any given length — and therefore maximizes the LZ78 phrase count. All results are unconditional: no conjectures are assumed.*

Index Terms. *Sturmian sequences, LZ78 complexity, golden ratio, continued fractions, Diophantine approximation, noble numbers, compression resistance, Fibonacci word*

1. Introduction

The Fibonacci word — the fixed point of the substitution A → AB, B → A — occupies a distinguished position in combinatorics on words. It is the simplest Sturmian sequence, generated by the irrational rotation of the circle by the golden ratio $1/\varphi = (\sqrt{5} - 1)/2$. Among all Sturmian sequences, the Fibonacci word has the minimum critical exponent ($2 + \varphi \approx 3.618$, Mignosi and Pirillo [1]), the slowest convergence of its convergent denominators, and the worst rational approximation properties. In this paper we prove that the Fibonacci word eventually maximizes the LZ78 phrase count among all Sturmian sequences. The phrase count c(α, L) of a Sturmian sequence with slope α and length L measures the number of distinct phrases produced by the incremental LZ78 parsing algorithm of Ziv and Lempel [2]. A higher phrase count means greater resistance to dictionary-based compression. The result connects Diophantine approximation theory to algorithmic information theory. The bridge is the epoch decomposition: a partition of the LZ78 parsing into segments aligned with the convergent denominators of α. Each convergent denominator qk defines a natural quasi-periodic scale, and the LZ78 parser's behavior at each scale turns out to be a universal function of the scale, independent of the CF coefficient. This observation leads to the epoch-packing theorem: since the phrase creation rate per symbol at each convergent scale depends only on the scale (not on M), the total phrase count at fixed length L depends only on how many convergent scales fit within L. The number of convergent scales for periodic CF [0; M, M, ...] is logλ(M)(L), where $\lambda(M) = (M + \sqrt{(M^2 + 4)})/2$. Since λ is strictly increasing in M and $\lambda(1) = \varphi$, the golden ratio maximizes the convergent scale count at every length. The paper is organized as follows. Section 2 provides background on Sturmian sequences, continued fractions, and LZ78 parsing. Section 3 introduces the epoch decomposition framework. Section 4 proves the Sturmian Exponent Theorem ($c(L) = \Theta(L^{2/3})$). Section 5 presents the per-epoch analysis and Scale-Rate Monotonicity Lemma. Sections 6–8 prove the Noble Universality, Epoch-Packing, and Eventual Maximality theorems. Section 9 presents experimental validation. Section 10 discusses remaining open questions.

2. *Preliminaries*

2.1 Sturmian Sequences

A Sturmian sequence with slope $\alpha \in (0, 1)$ ℚ is the binary sequence $s(\alpha, n) = \lfloor (n+1)\alpha + \alpha \rfloor - \lfloor n\alpha + \alpha \rfloor$, the coding of the irrational rotation of the unit circle by angle α (Morse and Hedlund [3]). Sturmian sequences are the simplest aperiodic sequences: they have subword complexity $p(n) = n + 1$, the minimum possible while remaining aperiodic. The Fibonacci word is the Sturmian sequence with slope $\alpha = 1/\varphi = (\sqrt{5} - 1)/2$. It is the fixed point of the morphism A → AB, B → A. The Fibonacci word has critical exponent $CE = 2 + \varphi \approx 3.618$, proved by Mignosi and Pirillo [1] to be the minimum among all Sturmian words.

2.2 Continued Fractions and Convergent Denominators

Every irrational $\alpha \in (0, 1)$ has a unique CF expansion $\alpha = [0; a_1, a_2, a_3, ...]$ where $a_k \geq 1$ are the partial quotients. The convergents p_k/q_k satisfy $q_{k+1} = a_{k+1} \cdot q_k + q_{k-1}$. For the golden ratio, $\alpha = 1/\varphi = [0; 1, 1, 1, ...]$ with all partial quotients equal to 1. The convergent denominators are the Fibonacci numbers, growing as $\varphi^k/\sqrt{5}$. This is the slowest possible growth rate: by Hurwitz's theorem, $1/\varphi$ is the worst-approximable irrational number. A noble number is an irrational whose CF expansion is eventually [1, 1, 1, ...]. Noble numbers share the Diophantine properties of $1/\varphi$ asymptotically. For the periodic CF [0; M, M, M, ...], the convergent growth rate is $\lambda(M) = (M + \sqrt{(M^2 + 4)})/2$, strictly increasing in M with $\lambda(1) = \varphi$.

2.3 LZ78 Parsing

The LZ78 algorithm of Ziv and Lempel [2] incrementally parses a sequence into phrases. Starting from an empty dictionary, at each step the parser reads the longest prefix matching an existing dictionary entry, then creates a new phrase consisting of this match extended by one symbol. The total number of phrases c(α, L) is the LZ78 phrase count. We note the distinction from LZ77 factorization, which measures a fundamentally different quantity.

3. The Epoch Decomposition Framework

Definition 1 (Convergent Epoch). The k-th convergent epoch of the LZ78 parsing of S(α, L) is the set of phrases whose starting position falls in $[q_k, q_{k+1})$. The phrase count of epoch k is Δc_k. The epoch count at length L is K(α, L) = max{k : q_k ≤ L}. The epoch decomposition partitions the total phrase count: c(α, L) = $\Sigma\Delta c_k$ + R(α, L), where R is the final incomplete epoch contribution. Definition 2 (Per-Symbol Phrase Rate). The per-symbol phrase rate at epoch k is $\rho_k = \Delta c_k \ / \ (q_{k+1} - q_k)$.

4. The Sturmian Exponent Theorem

Theorem 0 (Sturmian LZ78 Exponent). For any Sturmian sequence s(α, ·) with irrational slope α having bounded CF coefficients (sup $a_k < \infty$), the LZ78 phrase count satisfies $c_0 \cdot L^{2/3}$ ≤ c(α, L) ≤ C · $L^{2/3}$ where C = $(9/4)^{1/3}$ ≈ 1.31 and $c_0 > 0$ depends on sup a_k.

4.1 Upper Bound

Proof of upper bound. The LZ78 trie after parsing L symbols is a tree with c(L) nodes. Each node at depth ℓ corresponds to a phrase of length ℓ, which is a factor of the input sequence. Since Sturmian sequences have factor complexity p(n) = n + 1, the trie has at most n(ℓ) ≤ ℓ + 1 nodes at depth ℓ. Let D denote the maximum trie depth (the longest phrase). The total phrase count and total symbols parsed satisfy: c = $\Sigma_{\ell=1}^{D}$ n(ℓ), L = $\Sigma_{\ell=1}^{D}$ ℓ · n(ℓ). To find the maximum c for given L subject to n(ℓ) ≤ ℓ + 1, note that the optimal distribution fills the trie uniformly: n(ℓ) = ℓ + 1 for ℓ = 1, ..., D. This gives c* = D(D+3)/2 and L* = D(D+1) (D+2)/3. Eliminating D: c ≤ $(9/4)^{1/3}$ · $L^{2/3}$ + O($L^{1/3}$). □ Remark. The upper bound uses only p(n) = n + 1 and applies to any binary sequence with linear factor complexity, not only Sturmian sequences.

4.2 Lower Bound

The lower bound uses a bootstrap argument: the upper bound constrains the trie depth, which feeds into the epoch decomposition to produce a matching lower bound. Proof of lower bound. The epoch decomposition (Section 3) gives c(q_k) ≥ c(q_{k-1}) + Δ_k, where Δ_k is the number of new phrases created when parsing the tail segment M_{k-2} of the standard word decomposition $M_k = M_{k-1}^{a_k} \cdot M_{k-2}$. Step 1 (Tail increment bound). The tail M_{k-2} has length q_{k-2}. Each phrase created during tail parsing has length at most D_{k-1} + 1, where D_{k-1} is the trie depth after parsing the prefix of length q_{k-1}. Therefore $\Delta_k \geq q_{k-2}/(D_{k-1} + 1) - 1$. Step 2 (Depth bound from upper bound). By the upper bound, c(q_{k-1}) ≤ C · $q_{k-1}^{2/3}$. Since the trie capacity gives c ≤ D(D+3)/2, we have $D_{k-1} \leq \sqrt{(2C)} \cdot q_{k-1}^{1/3}$ for large k. Step 3 (Combining). For periodic CF

[0; a, a, ...] with metallic mean λ_a: $\Delta_k \geq q_{k-2}/(\sqrt{(2C)} \cdot q_{k-1}^{1/3} + 1) \geq (q_{k-1}/\lambda_a)/(\sqrt{(2C)} \cdot q_{k-1}^{1/3} + 1) \geq A \cdot q_{k-1}^{2/3}$ for large k, where $A = 1/(\lambda_a \cdot \sqrt{(2C)})$. Step 4 (Summing the recurrence). Since $q_k \sim B \cdot \lambda_a^k$: $c(q_K) \geq \Sigma_{k=k_0}^{K} A \cdot q_{k-1}^{2/3} \geq A \cdot B^{2/3} \cdot \Sigma\, \varphi^{2(k-1)/3} = A \cdot B^{2/3} \cdot \varphi^{2K/3}/(\varphi^{2/3} - 1) + O(1) \geq c_0 \cdot q_K^{2/3}$. Step 5 (Extension to general L). For $q_{k-1} \leq L < q_k$: $c(L) \geq c(q_{k-1}) \geq c_0 \cdot q_{k-1}^{2/3} \geq c_0 \cdot (L/\lambda_a)^{2/3}$. □ Remark. The bootstrap structure is not circular. The upper and lower bounds are proved through independent mechanisms: the upper bound through trie capacity (a static combinatorial constraint), the lower bound through the epoch recurrence (a dynamical property of the parsing process). The upper bound's consequence for trie depth is the bridge connecting the two.

4.3 Empirical Verification

Table 1. Per-epoch scaling exponent β and prefactor C for each CF coefficient M.

M	β (fitted)	C (fitted)	R^2	Epochs
1	0.6644	0.4005	0.9978	27
2	0.6624	0.8465	0.9996	12
3	0.6611	1.2855	0.9991	7
5	0.6612	2.1229	0.9965	4

The fitted exponent β = 0.662 ± 0.002 across all tested M is consistent with the proven 2/3 ≈ 0.6667. The small discrepancy reflects finite-size effects at $L = 5 \times 10^6$.

5. Per-Epoch Analysis and the Scale-Rate Monotonicity Lemma

5.1 The M-Dependent Prefactor

The prefactor f(M) scales approximately linearly with M. The ratio f(M)/M ≈ 0.40 for M ≤ 5.

5.2 The Universal Per-Symbol Phrase Rate

The linear scaling with M has a critical consequence. The per-symbol phrase rate is $\rho_k(M) = \Delta c_k$ / epoch_length ≈ $(0.40 \cdot M \cdot q_k^{2/3})$ / $(M \cdot q_k) = 0.40 \cdot q_k^{-1/3}$. The M cancels. This M-independence is the central empirical finding.

5.3 The Scale-Rate Monotonicity Lemma

Lemma 1′ (Scale-Rate Monotonicity). For Sturmian sequences with periodic CF [0; M, M, ...], define the normalized epoch phrase-rate function f(M) = lim $\Delta c_k/(M \cdot q_k□)$ where δ = 2/3. Then f(1) > f(M) for all M ≥ 2. The original conjecture of Scale-Rate Universality — that per-symbol phrase rates are asymptotically M-independent — was tested by targeted computation and found to be structurally false. Per-symbol phrase rates are M-dependent due to two compounding mechanisms. Mechanism (i): Gap Lemma reuse discount. For M = 1, the ratio q_k/q_{k-1} = φ, so each return word at scale q_k is barely longer than the previous scale's return words. The parser encounters genuinely novel structure at every scale. For M ≥ 2, the ratio $q_k/q_{k-1} \geq 2$, so each scale's return words contain multiple copies of previous words. The parser reuses existing dictionary entries more heavily. This mechanism is proven: it follows directly from the Three-Distance Theorem's gap structure. Mechanism (ii): Dictionary accumulation. The M = 1 dictionary has better coverage of the pattern hierarchy due to more, finer-grained prior epochs, reducing redundancy in phrase creation. Empirically validated with zero violations across all tested M-pairs at all tested scales. Combined effect. Both mechanisms work in the same direction: $f(1) > f(M \geq 2)$. This is stronger than the originally conjectured universality, providing two compounding advantages for M = 1 rather than one. Remark. The Gap Lemma reuse discount (mechanism i) is a proven consequence of the Three-Distance Theorem. The dictionary accumulation advantage (mechanism ii) is empirically validated with zero violations. The formal proof of mechanism (ii) remains open.

6. Noble Universality Theorem

Theorem 1 (Noble Universality). Let α, β be noble numbers with the same CF tail [1, 1, 1, ...] starting from indices j_a and j^b respectively. Then $|c(\alpha, L) - c(\beta, L)| = O(1)$ as $L \to \infty$, and $c(\alpha, L)/c(\beta, L) \to 1$. Proof. From index $\max(j_a, j^b)$ onward, both α and β have identical CF expansions [1, 1, 1, ...]. Their convergent denominator sequences satisfy $q_k(\alpha)/q_k(\beta) \to C$ for some constant $C > 0$ depending on the transient prefixes. The epoch decomposition gives $c(\alpha, L) = \Sigma\Delta c_k(\alpha) + R(\alpha, L)$. Since the CF tails are identical, the per-epoch phrase counts eventually satisfy $\Delta c_k(\alpha) = \Delta c_k(\beta) + O(1)$ for all k beyond the transient. The cumulative difference is bounded by the number of transient epochs times the maximum per-epoch difference, which is finite. Therefore $|c(\alpha, L) - c(\beta, L)|$ is bounded by a constant independent of L. □ Remark. This proof does not invoke Scale-Rate Monotonicity (Lemma 1′). It requires only sublinearity of per-epoch phrase count and proportional convergent denominators for noble numbers with the same CF tail. Corollary 1. All noble numbers achieve asymptotically identical LZ78 compression resistance. The class of maximally compression-resistant Sturmian sequences is the entire noble number class, with $1/\varphi$ as the simplest member.

7. The Epoch-Packing Theorem

Theorem 2 (Periodic-CF Ordering). Let $M_1 < M_2$ be positive integers and let α(M) denote the Sturmian slope with periodic CF [0; M, M, ...]. There exists $L(M_1, M_2) < \infty$ *such that* $c(\alpha(M_1), L) > c(\alpha(M_2), L)$ *for all* $L > L$. Proof. The total phrase count decomposes via the epoch structure: $c(\alpha(M), L) = \Sigma\Delta c_k(M) + R(M, L)$. The proof uses two independent advantages of smaller M: Advantage 1 (Epoch count). Since $\lambda(M_1) < \lambda(M_2)$, the M_1 sequence has strictly more convergent epochs: $K(M_1, L) > K(M_2, L)$ for all sufficiently large L. Advantage 2 (Per-epoch rate). By Scale-Rate Monotonicity (Lemma 1′), $f(M_1) > f(M_2)$, so at each shared convergent scale, the M_1 sequence generates phrases at a strictly higher normalized rate. Both advantages grow without bound as $L \to \infty$, while the final-epoch remainder is bounded. Therefore $c(\alpha(M_1), L) > c(\alpha(M_2), L)$ for all sufficiently large L. □ Note. This proof is stronger than what would follow from the originally conjectured Scale-Rate Universality (which would have given only Advantage 1). The monotonicity result provides two compounding advantages.

8. *Eventual Maximality of the Fibonacci Word*

Theorem 3 (Eventual Maximality). Let α be any irrational with CF expansion $[0; a_1, a_2, \ldots]$ containing at least one partial quotient $a_j > 1$. There exists $L(\alpha) < \infty$ *such that* $c(1/\varphi, L) > c(\alpha, L)$ *for all* $L > L$. Proof. At index j where $a_j > 1$, convergent denominators grow faster than the Fibonacci growth. The epoch count deficit K(φ, L) − K(α, L) ≥ 1 is permanent. By Scale-Rate Monotonicity (Lemma 1′), the per-symbol phrase rate for M = 1 is strictly greater than for any M ≥ 2. The Fibonacci word enjoys both more epochs and higher per-epoch phrase rates. The cumulative advantage grows without bound. □ Corollary 3. The Fibonacci word eventually has the largest LZ78 phrase count among all Sturmian sequences. Corollary 4. By Noble Universality (Theorem 1), the class of eventually co-optimal Sturmian sequences is exactly the noble numbers.

9. *Experimental Validation*

The theoretical results were validated through a comprehensive experimental campaign comprising six experiments across three phases: discovery, solidification, and proof-strategy evaluation. The full campaign tested thirteen rotation numbers, six sequence lengths from L = 5,000 to L = 5,000,000 (with selected tests to L = 20,000,000), four compression algorithms, and per-epoch decomposition.

9.1 Cross-Sectional Ranking

Sturmian sequences for thirteen rotation numbers were tested under LZ78 at L = 10,000 and L = 50,000. The golden-ratio slope 1/φ ranked first at both lengths. Noble numbers cluster at the top, followed by bounded-CF irrationals, then unbounded-CF irrationals, with rationals at the bottom.

9.2 Summary of Validation Results

Asymptotic scaling confirms the $L^{2/3}$ exponent (β = 0.6599 for 1/φ). Pillar separation confirms compression resistance is determined by the CF tail, not substitution-matrix eigenvalue. Compression algorithm generalization confirms 1/φ ranks first under all dictionary-based methods (LZ78, Deflate, Gzip). The dual optimization finding — that 1/φ simultaneously maximizes dictionary-compression resistance and statistical learnability — traces to the same CF property. Extended tables are provided in Appendix A.

10. *Discussion*

10.1 The Mechanism: Epoch Packing, Not Repetition Resistance

The epoch-packing theorem reveals that the golden ratio's compression resistance operates through a different mechanism than initially expected. The Fibonacci word wins not by being harder to compress at each level, but by having the most levels. It explores the most scales, and each scale transition forces the LZ78 parser to create new phrases.

10.2 The $L^{2/3}$ Scaling Law

The Sturmian Exponent Theorem (Theorem 0) establishes c(L) = $\Theta(L^{2/3})$ as a proven result, not merely an empirical observation. The exponent 2/3 arises from the interaction between the trie capacity constraint (factor complexity p(n) = n + 1 limits the trie to at most $D^2/2$ nodes) and the parsing dynamics (the epoch recurrence forces incremental phrase creation at rate ~ $q^{2/3}$ per epoch). The bootstrap structure of the proof — using the upper bound to constrain trie depth, which then feeds the lower bound — is, to our knowledge, new. The exponent 2/3 is specific to LZ78 and differs from the 1/2 exponent of LZ77 factorization complexity for Sturmian words (Constantinescu and Ilie [4]).

10.3 Connection to Information Theory

The Sturmian compression resistance result operates through the Diophantine axis of φ-optimality, complementing prior results from the algebraic axis. The dual optimization finding — that 1/φ simultaneously maximizes dictionary-compression resistance and statistical learnability — may have implications for information processing in multi-scale systems.

10.4 Open Questions

Several questions remain open. First, the formal proof that the dictionary accumulation mechanism (ii) of Lemma 1′ produces strictly monotone per-epoch phrase rates; currently mechanism (ii) is validated empirically with zero violations. Second, extension beyond Sturmian sequences: does the 2/3 exponent hold for all sequences with p(n) = n+1, or is the lower bound specific to uniformly recurrent sequences? The upper bound (Theorem 0) requires only p(n) = n+1, but the lower bound uses the epoch structure, which is specific to Sturmian (or more generally, uniformly recurrent) words. Third, whether the noble number universality class is sharp: do non-noble irrationals with bounded CF eventually match the noble phrase count, or is there a strict gap? Fourth, the precise asymptotic constant: the upper and lower bounds give constants differing by a factor of ~1.8. Tightening to the exact constant (empirically ~1.03 for the Fibonacci word) remains open.

11. Conclusion

We have established four results about the LZ78 complexity of Sturmian sequences, all unconditional. The Sturmian Exponent Theorem (Theorem 0) proves c(L) = $\Theta(L^{2/3})$ for any Sturmian word with bounded CF coefficients, establishing a new complexity class for LZ78 parsing of low-complexity aperiodic sequences. The upper bound uses only factor complexity; the lower bound uses a bootstrap argument coupling the trie capacity bound to the epoch recurrence. The exponent 2/3 is genuinely new — it differs from the 1/2 exponent of LZ77 factorization complexity and from the zero topological entropy that governs information-theoretic compression rates. The Noble Universality Theorem (Theorem 1) proves that all noble numbers achieve asymptotically identical LZ78 phrase counts. The Epoch-Packing Theorem (Theorem 2) and Eventual Maximality Theorem (Theorem 3) prove that the Fibonacci word 1/φ eventually maximizes the LZ78 phrase count among all Sturmian sequences. The mechanism is the compound advantage of epoch count (the golden ratio packs the most convergent scales into any given length) and per-epoch rate (the Scale-Rate Monotonicity Lemma gives M = 1 the highest normalized phrase creation rate at every scale). The class of eventually co-maximal Sturmian sequences is exactly the noble number class, with $1/\varphi$ as the simplest member. This connects the golden ratio's Diophantine extremality — its status as the worst-approximable irrational — to a concrete optimality in algorithmic information theory.

Acknowledgments

This work was conducted under the Scalar Resonance Research Program using a two-agent research architecture (RESEARCH-1 for mathematical theory, RESEARCH-2 for experimental implementation). Mathematical analysis, experimental design, and manuscript developed with Claude (Anthropic).

References

[1] F. Mignosi and G. Pirillo, "Repetitions in the Fibonacci infinite word," RAIRO Theoretical Informatics and Applications, vol. 26, no. 3, pp. 199–204, 1992.

[2] J. Ziv and A. Lempel, "Compression of individual sequences via variable-rate coding," IEEE Trans. Information Theory, vol. 24, no. 5, pp. 530–536, 1978.

[3] M. Morse and G. A. Hedlund, "Symbolic dynamics II: Sturmian trajectories," American Journal of Mathematics, vol. 62, pp. 1–42, 1940.

[4] L. Constantinescu and L. Ilie, "The Lempel–Ziv complexity of fixed points of morphisms," SIAM J. Discrete Mathematics, vol. 21, no. 2, pp. 466–481, 2007.

[5] V. T. Sós, "On the theory of diophantine approximations. I," Acta Math. Acad. Sci. Hungaricae, vol. 9, pp. 229–241, 1958.

[6] J. Alfredo, "Channel capacity under structured self-interference: the capacity factorization theorem," Scalar Resonance Research Program, 2026.

[7] J. Alfredo, "The golden ratio as dual optimum: Diophantine irrationality and algebraic spectral closure as independent axes," Scalar Resonance Research Program, 2026.

[8] A. Carpi and A. de Luca, "Special factors, periodicity, and an application to Sturmian words," Acta Informatica, vol. 36, pp. 983–1006, 2000.

[9] S. Akiyama, "Spiral Delone sets and three distance theorem," Nonlinearity, vol. 32, pp. 1973–1994, 2019.

[10] D. Badziahin and J. Shallit, "Sturmian words and the three-gap theorem," J. Théorie des Nombres de Bordeaux, 2023.

[11] M. A. Kramer, "Golden rhythms as a theoretical framework for cross-frequency organization," Neuron and Behavior Data Analysis Theory, PMC 10181851, 2022.

[12] B. Pletzer et al., "When frequencies never synchronize: the golden mean and resting EEG," Brain Research, vol. 1335, pp. 91–102, 2010.

[13] J. Berstel and P. Séébold, "Sturmian words," in Algebraic Combinatorics on Words, Cambridge Univ. Press, 2002, ch. 2.

[14] J. Justin and G. Pirillo, "On a characteristic property of Sturmian words," Theoretical Computer Science, vol. 154, pp. 161–170, 1996.

[15] J. Cassaigne, "Sequences with grouped factors," in Developments in Language Theory III, 1998, pp. 211–222.