Rigidity of Pattern-Avoiding Breadth-First Reading Words of Increasing Trees

Victor Bona
Independent Researcher, Blumenau, Brazil
victor.bona@hotmail.com

July 2026

Abstract

An increasing tree on n nodes carries the labels $1, \ldots, n$ so that every child exceeds its parent, and its breadth-first reading word is the permutation obtained by reading labels level by level from left to right. Levin, Pudwell, Riehl, and Sandberg introduced the enumeration of permutations that avoid a classical pattern of length three and arise as such reading words; six OEIS sequences record the counts for unary-binary and full binary trees against the patterns 231, 312, and 321, and all six had remained unextended since 2014. We prove two rigidity theorems. First, a 312-avoiding permutation is the breadth-first reading word of some increasing unary-binary tree if and only if it satisfies the binary-heap inequality $w_{\lfloor i/2 \rfloor} < w_i$; realizability over the infinite family of admissible shapes collapses to a single shape, and the resulting equality of sequences `A245899` and `A246747`, conjectured in the OEIS, holds with the identity map as the bijection. Second, every 231-avoiding reading word of odd length realizable on an increasing unary-binary tree is realizable on an increasing full binary tree, which proves `A245901`$(k) =$ `A245898`$(2k-1)$. Both collapses fail for the pattern 321, and we exhibit the minimal counterexamples. The central technical ingredient is a promotion cascade lemma that holds for all words with no pattern hypothesis: enlarging a breadth-first level by the next element never destroys realizability of the remaining suffix. Pattern avoidance enters at exactly one parity-driven step, which is precisely where 321 breaks. We also extend the six sequences by 28 new terms. All results are verified three ways: by self-contained proofs, by exhaustive computation, and by complete machine-checked formalizations in Lean 4 whose statements quantify over genuine inductive trees.

1 Introduction

Pattern avoidance began with Knuth's analysis of stack-sortable permutations [6] and has since become a central topic in enumerative combinatorics [3, 5]. A recurring modern theme replaces the ambient set S_n with the image of a combinatorial family under a reading map: one fixes a class of labeled structures, reads each structure into a permutation, and asks how many of the resulting permutations avoid a given pattern. Levin, Pudwell, Riehl, and Sandberg carried out this program for k-ary heaps read in breadth-first order [7], and Defant later settled their growth-rate conjecture for 321-avoiding binary heaps [4]. Related reading-word problems appear for linear extensions of forests and rectangular posets [1, 4] and for increasing unary-binary trees with an associated permutation avoiding 213, which are in bijection with basketball walks [2].

This paper concerns the breadth-first reading words themselves. An *increasing tree* on n nodes is a rooted plane tree whose nodes carry the labels $1, \ldots, n$, each child labeled with a larger value than its parent. Reading the labels level by level, left to right, produces the *breadth-first reading word*, a permutation of $\{1, \dots, n\}$ that necessarily begins with 1. Following [7], six OEIS [9] sequences count the permutations that avoid a single classical pattern of length three and occur as reading words: `A245898`, `A245899`, and `A245900` for increasing unary-binary trees (every node has at most two children) against the patterns 231, 312, and 321, and `A245901`, `A245902`, and `A245903` for increasing full binary trees (every node has zero or two children), whose words have odd length. All six entries carried the keyword `more`, requesting extension, and none had gained a term since their creation in 2014. Two further statements sat in the margins of the database as recorded observations rather than theorems. The entry `A246747`, which counts binary heaps on n elements whose reading word avoids 312, carried the note that it may equal `A245899`, and the data of the full binary sequences agreed with the odd-indexed terms of their unary-binary counterparts as far as either had been computed.

We resolve these questions completely for the patterns 231 and 312, and we show that both phenomena genuinely depend on the pattern by exhibiting explicit counterexamples for 321.

Our first theorem is a rigidity statement. Among all increasing unary-binary trees on n nodes there are exponentially many shapes, and a priori a permutation could be realizable as a reading word on some exotic shape only. For 312-avoiding permutations this freedom is illusory.

Theorem 1.1 (Heap collapse; Theorem 3.2 below). *A 312-avoiding permutation w of $\{1, \dots, n\}$ is the breadth-first reading word of an increasing unary-binary tree if and only if $w_{\lfloor i/2 \rfloor} < w_i$ for all $2 \le i \le n$, that is, if and only if w is the array of a binary heap.*

Since a heap on a fixed shape is determined by its reading word, the theorem identifies the 312-avoiding reading words of unary-binary trees with the 312-avoiding heaps of [7]: the conjectured equality `A245899`$(n) =$ `A246747`(n) holds for all n, and the bijection realizing it is the identity map on words. The complete shape on an odd number of nodes has no node with exactly one child, so the same theorem yields `A245902`$(k) =$ `A245899`$(2k-1)$: for 312-avoiding words of odd length, the unary-binary, full binary, and heap realizability notions coincide. Through the recurrence of `A246747`, the sequence `A245899` acquires the closed description

$$a(n) \;=\; \sum_{i=0}^{\lfloor (n-1)/2 \rfloor} C_i \, a(n-i-1), \qquad a(0) = 1,$$

with C_i the i-th Catalan number.

Our second theorem is a parity statement for the pattern 231, where no single-shape collapse holds.

Theorem 1.2 (Parity collapse; Theorem 4.5 below). *Every 231-avoiding permutation of odd length that is the breadth-first reading word of an increasing unary-binary tree is also the breadth-first reading word of an increasing full binary tree. Consequently `A245901`$(k) =$ `A245898`$(2k-1)$ for all k.*

Both theorems are sharp in the pattern. For 321 the corresponding identity fails at the first length beyond the previously published data: among 321-avoiding words of length 11 there are 8095 unary-binary reading words and only 8048 full binary reading words, and the 47 missing words are listed explicitly in the repository accompanying this paper.

The proofs share a two-part architecture that we believe is of independent interest. Realizability of a word is a property of its partition into breadth-first levels together with an order-preserving assignment of children to parents. The first part of each proof canonicalizes the assignment within a fixed partition: for 312 the balanced assignment always works (Lemma 3.1), while for 231 an odd level can always absorb the next element of the word (Lemma 4.3). The second part reshapes the partition itself, and here the key tool is a lemma with no pattern hypothesis at all.

Lemma 1.3 (Promotion cascade; Lemma 4.1 below)**.** *Let w be any word. If the suffix of w beginning with an element x can be arranged into breadth-first levels below a previous level P, then the suffix with x removed can be arranged into levels below P extended by x, and the total depth strictly decreases.*

The cascade moves x up into the parent level and lets the old children of x ride along one level upward, each promoted node keeping its original parent. No new comparison between labels is ever created, which is why no pattern condition is needed. Pattern avoidance is consumed at exactly one point in the whole development, the parity step of Lemma 4.3, and the failure of that single step for 321 is what makes the counterexamples possible.

Every result in this paper is verified three ways. The proofs below are self-contained. An exhaustive computation checks every lemma on all instances at small sizes, some three hundred thousand cases in total, and reproduces all thirty-six previously published terms of the six sequences by two independent algorithms. Finally, both theorems are formalized in Lean 4 [8]; the formal statements quantify over an inductive type of genuine rooted labeled trees, the development contains no unproven assumptions, and the kernel certifies both theorems from Lean's standard axioms. The formalization effort also simplified the informal proofs, and the versions printed here incorporate those simplifications. Code, data, certificates, and the Lean development are available in the accompanying repository.[1]

The paper is organized as follows. Section 2 fixes definitions and the level-matching formulation of realizability. Section 3 proves the heap collapse. Section 4 develops the promotion cascade and proves the parity collapse. Section 5 discusses sharpness at 321 and reports the new terms. Section 6 describes the computational and formal verification. Section 8 collects open problems.

2 Preliminaries

We write $[n] = \{1, \dots, n\}$ and use one-line notation $w = w_1 w_2 \cdots w_n$ for permutations. A permutation w *contains* the pattern $\tau \in S_3$ if some subsequence of w is order-isomorphic to τ, and *avoids* τ otherwise. We write $\mathrm{Av}_n(\tau)$ for the set of τ-avoiding permutations of $[n]$. Three instances of containment recur throughout, and we spell them out once: an occurrence of 231 is a triple of positions $i < j < l$ with $w_l < w_i < w_j$, an occurrence of 312 is a triple with $w_j < w_l < w_i$, and an occurrence of 321 is a decreasing triple $w_i > w_j > w_l$.

All trees are rooted and plane, meaning that the children of every node are linearly ordered. An *increasing tree* on n nodes is such a tree with labels $1, \dots, n$, bijectively assigned, in which every child's label exceeds its parent's label; the root is therefore labeled 1. A tree is *unary-binary* if every node has at most two children and *full binary* if every node has zero or two children. The *breadth-first reading word* $\mathrm{bfs}(T)$ lists the labels of T level by level, left to right, starting with the root. Figure 1 shows an example. We write $\mathrm{UB}_n(\tau)$ for the set of words in $\mathrm{Av}_n(\tau)$ realizable on increasing unary-binary trees and $\mathrm{FB}_n(\tau)$ for those realizable on increasing full binary trees. A full binary tree on n nodes exists only for odd n, and since every full binary tree is unary-binary we have $\mathrm{FB}_n(\tau) \subseteq \mathrm{UB}_n(\tau)$.

[1] `https://github.com/vicotrbb/bfs-avoidance-frontier`

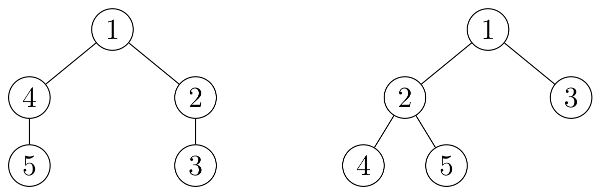

Figure 1: Two increasing unary-binary trees. The left tree has reading word 14253; the right tree is the complete (heap) shape on five nodes and has reading word 12345. The word 14253 contains 312 (as the subsequence 4 2 3) and indeed violates the heap inequality at position 5.

Realizability is conveniently phrased without trees. A *level partition* of w cuts its positions into consecutive nonempty blocks $L_0, L_1, \ldots, L_m$ with L_0 the single first position. A *matching* between an adjacent pair (P, C), where $P = (p_1, \ldots, p_M)$ lists the values of a block and $C = (c_1, \ldots, c_K)$ those of the next block, is a partition of C into consecutive, possibly empty groups $g_1, \ldots, g_M$, one per parent in order, such that every group has size at most two and every member of g_t exceeds p_t. For full binary trees each group must have size zero or two. Unwinding the definitions gives the following equivalence, which we record for later use; in the Lean development it is a proved theorem connecting the matching formulation to an inductive type of trees.

Proposition 2.1. *A permutation w is the breadth-first reading word of an increasing unary-binary tree if and only if w admits a level partition with a valid matching between every adjacent pair of blocks, and likewise for full binary trees with groups of size zero or two.*

Two closure properties of matchings are immediate and will be used silently: removing the first element of the child block preserves validity (its group shrinks), and appending a new parent at the end of the parent block together with a final group of new children preserves validity. Summing group sizes gives the capacity bound $|C| \leq 2|P|$.

Finally, since blocks of a level partition are contiguous in w, any elements drawn from distinct blocks appear in w in block order, so pattern conditions on w restrict the relative values across blocks. This is the only mechanism through which avoidance enters the proofs.

3 The heap collapse for 312

Throughout this section w avoids 312. The direction from heaps to trees is immediate: the complete shape, in which position i has children at positions $2i$ and $2i + 1$, is unary-binary, and the heap inequality makes its labeling increasing. The substance is the converse, and it rests on a canonicalization of matchings.

Lemma 3.1 (Balanced matching)**.** *Let (P, C) be an adjacent pair of blocks of a 312-avoiding word admitting a valid matching. Then the balanced matching, which assigns child c_j to parent $p_{\lceil j/2 \rceil}$, is valid: $c_j > p_{\lceil j/2 \rceil}$ for all j.*

Proof. Fix a valid matching φ and write $t = \lceil j/2 \rceil$. Groups have size at most two, so the first j children occupy at least t parents and $\varphi(j) \geq t$. If $\varphi(j) = t$ then $c_j > p_t$ directly. Otherwise $\varphi(j) > t$ and $c_j > p_{\varphi(j)}$. Suppose $c_j < p_t$. Then $p_{\varphi(j)} < c_j < p_t$, while the positions of p_t, $p_{\varphi(j)}$, c_j in w are increasing. This triple is an occurrence of 312, a contradiction. Since labels are distinct, $c_j > p_t$. $\square$

Theorem 3.2 (Heap collapse)**.** *Let $w \in \mathrm{Av}_n(312)$. Then w is the breadth-first reading word of an increasing unary-binary tree if and only if $w_{\lfloor i/2 \rfloor} < w_i$ for all $2 \le i \le n$.*

Proof. Only the forward direction remains. Fix a realization of w with level sizes $n_0 = 1, n_1, \ldots, n_m$ and, by Lemma 3.1, balanced matchings throughout. Let $s_j = 1 + \sum_{t<j} n_t$ denote the 1-indexed starting position of level j.

We first claim $s_j \le 2s_{j-1}$ for all $j \ge 1$. Equivalently, writing $d_j = s_{j+1} - 2s_j = n_j - s_j$, we must show $d_{j-1} \le 0$. The capacity bound gives $n_j \le 2n_{j-1}$, and therefore

$$d_j \;=\; n_j - s_j \;\le\; 2n_{j-1} - (s_{j-1} + n_{j-1}) \;=\; n_{j-1} - s_{j-1} \;=\; d_{j-1},$$

so the quantity d_j is nonincreasing, and $d_0 = n_0 - s_0 = 1 - 1 = 0$.

Now fix a position $p \ge 2$, lying in level j at local offset $k = p - s_j \ge 0$, and let $q = \lfloor p/2 \rfloor$ be its heap parent. Under the balanced matching the tree parent of p sits at position $b = s_{j-1} + \lfloor k/2 \rfloor$, and $w_b < w_p$. Monotonicity of the floor together with $s_j \le 2s_{j-1}$ gives

$$q \;=\; \left\lfloor \tfrac{s_j + k}{2} \right\rfloor \;\le\; \left\lfloor \tfrac{2s_{j-1} + k}{2} \right\rfloor \;=\; s_{j-1} + \left\lfloor \tfrac{k}{2} \right\rfloor \;=\; b,$$

and clearly $b < p$. If $q = b$ we are done. If $q < b$, suppose $w_q > w_p$. Then the values at the increasing positions $q < b < p$ satisfy $w_b < w_p < w_q$, an occurrence of 312. Hence $w_q < w_p$, as labels are distinct. □

Corollary 3.3. *For all n and k:*

(i) *A245899*(n) = *A246747*(n)*, and the underlying sets of words are equal, so the bijection between* 312*-avoiding unary-binary reading words and* 312*-avoiding binary heaps is the identity map.*

(ii) *A245902*(k) = *A245899*$(2k-1)$*, and for odd length the three word classes, unary-binary, full binary, and heap, coincide.*

(iii) *A245899 satisfies* $a(n) = \sum_{i \le (n-1)/2} C_i \, a(n-i-1)$.

Proof. On a fixed shape a heap is determined by its reading word, so heaps biject with heap words and (i) follows from Theorem 3.2. For (ii), a complete shape on an odd number of nodes has no node with exactly one child, hence is full binary; combined with FB $\subseteq$ UB and Theorem 3.2 the three classes coincide. Item (iii) transports the recurrence of A246747, established on the heap side following [7], across the equality (i). □

Remark 3.4. Theorem 3.2 is a rigidity phenomenon: the class of admissible shapes is exponentially large, yet for 312-avoiding words every realizable word is realizable on the one canonical shape. The proof shows slightly more, namely that any realization can be pushed onto the complete shape level by level without changing the word.

4 The parity collapse for 231

The pattern 231 admits no analogue of Theorem 3.2: the counts of A245898 differ from the 231-avoiding heap counts already at $n = 4$. What survives is a parity statement, and its proof requires reshaping level partitions rather than collapsing them onto a fixed shape. We first develop the reshaping tool in full generality.

4.1 The promotion cascade

For a level partition presented as a list of levels $\mathcal{L} = (L_1, \ldots, L_m)$ below a previous block P, define the *total depth* $\delta(\mathcal{L}) = \sum_{t=1}^{m} t \cdot |L_t|$.

Lemma 4.1 (Promotion cascade)**.** *Let w be an arbitrary word, let P be a block, and suppose the suffix $xu_1u_2 \cdots u_r$ of w admits a level partition $\mathcal{L}$ with valid matchings below P. Then the suffix $u_1 \cdots u_r$ admits a level partition $\mathcal{L}'$ with valid matchings below the extended block Px, with $\delta(\mathcal{L}') < \delta(\mathcal{L})$.*

Proof. The element x is the first element of the first level L_1. Set $J_0 = \{x\}$ and, inductively, let $J_t \subseteq L_{t+1}$ be the set of children, under the given matchings, of the nodes in J_{t-1}. Because matchings are order-preserving and each J_{t-1} is an initial segment of its level consisting of the earliest parents that still have children, each J_t is an initial segment of L_{t+1}; the sets J_t are the *promoted* nodes. Define new levels

$$L'_t \;=\; \bigl(L_t \setminus J_{t-1}\bigr) \cdot J_t,$$

that is, the old level with its promoted head removed and the next promoted set appended at its end. Positions within w remain contiguous: the block boundaries simply shift right by $|J_{t-1}|$ on the left and $|J_t|$ on the right.

The matchings for the new partition are assembled from old parent-child pairs only. Children in $L_t \setminus J_{t-1}$ keep their old parents, which lie in $L_{t-1} \setminus J_{t-2}$ precisely because the children of J_{t-2} are exactly J_{t-1}. Children in J_t also keep their old parents, namely the nodes of J_{t-1}, which now sit at the end of L'_{t-1}; group sizes and the order-preserving property are inherited. At the top, the elements of J_1, the old children of x, attach to x itself, which is now the last element of the extended block Px, while the rest of L'_1 keeps its old parents in P. No comparison between labels occurs anywhere in this construction beyond those already present in the old matchings, which is why no pattern hypothesis is needed. If some L'_t is empty, then all deeper new levels are empty as well, since a nonempty L_{t+1} forces $J_t \neq \emptyset$, and the partition simply shortens.

Finally, every promoted node moves up one level and no node moves down, while x leaves the level system entirely; hence δ decreases by at least one. □

Lemma 4.1 may be read as a strong exchange property of breadth-first realizability: the set of feasible previous-block extensions is closed under absorbing the next element of the word, for every word. In the formal development this lemma is proved exactly as stated, with no side conditions.

4.2 Two parity lemmas for 231

Avoidance of 231 enters through the following two lemmas. Both are proved by induction on $|P| + |C|$, and both consume the pattern through a single forced inequality.

Lemma 4.2 (Suffix compatibility)**.** *Let (P, C) be blocks of a* 231*-avoiding word, in block order. If $c_j > p_t$ and $j' > j$, then $c_{j'} > p_t$.*

Proof. Otherwise $(p_t, c_j, c_{j'})$ occupies increasing positions with $c_{j'} < p_t < c_j$, an occurrence of 231. □

Lemma 4.3 (Odd extension)**.** *Let (P, C) admit a valid matching in a* 231*-avoiding word, let $|C|$ be odd, and let x be the element of w immediately following the block C. Then $(P,\, C \cdot x)$ admits a valid matching.*

Proof. Induct on $|P| + |C|$. Write $P = pP'$ and let g be the group of p in a valid matching, so $C = g \cdot C'$ with a valid matching between P' and C'.

If $|g|$ is even, then $|C'|$ is odd; the inductive hypothesis extends (P', C') by x, and reattaching the group g to p proves the claim.

If $|g| = 1$, say $g = (c)$ with $c > p$, there are two cases. If C' is empty, then $C = (c)$ and it suffices to give p the group (c, x), for which we need $x > p$: otherwise (p, c, x) occupies increasing positions with $x < p < c$, an occurrence of 231. If C' is nonempty, let c' be its first element. The parent of c' lies in P', so c' exceeds some later parent, and moreover $c' > p$: otherwise (p, c, c') is again an occurrence of 231. Regroup by giving p the group (c, c'); the remaining children C'', namely C' without its head, satisfy $|C''| = |C'| - 1$, which is odd since $|C| = 1 + |C'|$ is odd, and the matching between P' and C' restricts to one between P' and C''. The inductive hypothesis extends (P', C'') by x, and assembling the pieces proves the claim. □

Lemma 4.4 (Even regrouping)**.** *Let (P, C) admit a valid matching in a* 231*-avoiding word with $|C|$ even. Then (P, C) admits a valid matching in which every group has size zero or two.*

Proof. The induction of Lemma 4.3 applies verbatim with the parities exchanged. In the case $|g| = 1$ the block C' has odd size, hence is nonempty, and the same forced inequality $c' > p$ permits the regrouping (c, c'); the remaining children form an even block and the induction proceeds. The base cases produce only groups of size zero or two. □

4.3 Proof of the parity collapse

Theorem 4.5 (Parity collapse)**.** *Let $w \in \mathrm{Av}_n(231)$ with n odd. If w is the breadth-first reading word of an increasing unary-binary tree, then w is the breadth-first reading word of an increasing full binary tree. Consequently,* A245901$(k) =$ A245898$(2k-1)$ *for all $k \geq 1$.*

Proof. Among all level partitions of w with valid matchings, choose one, $\mathcal{P}\colon L_0, L_1, \ldots, L_m$, whose size vector $(n_1, \ldots, n_m)$ is lexicographically maximal; the set of realizations is finite and nonempty, so the choice is possible.

We claim every n_i with $i \geq 1$ is even. Suppose not, and let i be minimal with n_i odd. Since $\sum_{t \geq 1} n_t = n - 1$ is even, the number of odd levels is even, so level i is not the last, and the element x of w immediately after L_i exists; it is the first element of L_{i+1}. By Lemma 4.3 the pair $(L_{i-1},\, L_i \cdot x)$ admits a valid matching, and by Lemma 4.1, applied with previous block L_i and the suffix beginning at x, the remainder of the word admits a level partition with valid matchings below $L_i \cdot x$. Splicing these below the unchanged levels $L_0, \ldots, L_{i-1}$ produces a realization whose size vector agrees with that of $\mathcal{P}$ before index i and exceeds it at index i. This contradicts lexicographic maximality.

All levels of $\mathcal{P}$ are therefore even, and Lemma 4.4 converts its matchings, level by level, into matchings with all groups of size zero or two. The resulting structure is an increasing full binary tree with reading word w. The sequence identity follows since $\mathrm{FB}_{2k-1}(231) \subseteq \mathrm{UB}_{2k-1}(231)$ holds trivially and the reverse inclusion was just proved. □

Remark 4.6. The proof is effective. Iterating the extension step terminates because the lexicographic order on bounded size vectors is well founded, and in the formal development the same argument runs on the strictly decreasing total depth of Lemma 4.1. Extracting the algorithm gives a procedure that converts any unary-binary realization of an odd-length 231-avoiding word into a full binary one.

5 Sharpness and new data

Both theorems fail for the pattern 321, and they fail at the first opportunity beyond the previously published data.

Proposition 5.1. $\mathrm{UB}_{11}(321)$ *has* 8095 *elements while* $\mathrm{FB}_{11}(321)$ *has* 8048*. In particular the analogue of Theorem 4.5 for* 321 *is false, and so is any heap collapse in the sense of Theorem 3.2. The lexicographically least word realizable on an increasing unary-binary tree but on no increasing full binary tree is*

$$1\,4\,2\,5\,3\,7\,8\,9\,10\,6\,11.$$

The 47 witnesses were found and confirmed by two independent programs, and each witness is a finite certificate that can be checked by hand or machine. The structural reason for the failure is visible in the proofs: the forced inequalities of Lemmas 4.2 and 4.3 arise from triples forming the patterns 231 or 312, which avoidance of 321 does not exclude. The published data of the six sequences, in which the full binary counts agreed with the odd-indexed unary-binary counts for all three patterns as far as they had been computed, was in this respect misleading: the agreement for 321 is a small-size coincidence that ends exactly where the 2014 computations stopped.

Table 1 reports the new terms of all six sequences. Terms up to length 11 were computed independently by a brute-force enumeration of trees and by a pruned depth-first search over words with an incremental realizability test; the two programs agree with each other and with all thirty-six previously published values. The deeper terms come from the second program alone.

Unary-binary, length n					Full binary, length ℓ			
n	`A245898` (231)	`A245899` (312)	`A245900` (321)		ℓ	`A245901` (231)	`A245902` (312)	`A245903` (321)
9	667	222	753		11	6788	1601	8048
10	2098	544	2439		13	74969	12416	92716
11	6788	1601	8095		15	877865	105769	1125054
12	22349	4095	27350					
13	74969	12416	93885					
14	254848	33785	326396					

Table 1: New terms of the six sequences. The previously published data ended at $n = 8$ and $\ell = 9$. The identities $\texttt{A245901}(k) = \texttt{A245898}(2k-1)$ and $\texttt{A245902}(k) = \texttt{A245899}(2k-1)$, visible in the table, are Theorems 4.5 and 3.2; the divergence $8095 \neq 8048$ in the 321 columns is Proposition 5.1.

6 Verification

The claims of this paper are supported by three independent layers, each sufficient on its own to check the results, and all contained in the accompanying repository.

Exhaustive computation. Every lemma was tested on every instance within exhaustive small ranges before being proved: Lemma 4.3 on 27,364 instances, Lemma 4.4 on 18,677, and Lemma 4.1, in its unconditional form over all words including those avoiding nothing, on 127,127 instances, with no failures. The two enumeration programs are algorithmically unrelated, one enumerating labeled trees and collecting reading words into a set, the other enumerating pattern-avoiding words and deciding realizability by dynamic programming, so agreement between them on every overlapping value is meaningful evidence of correctness independent of the proofs. This discipline earned its keep during the research: three natural candidate lemmas, among them a mirrored version of Lemma 3.1 for 231, were falsified by counterexamples at $n \leq 9$ before the correct statements were found. The final proofs use none of the falsified statements.

Formalization. Both theorems are formalized in Lean 4 [8] in about two thousand lines. The development defines an inductive type of rooted plane labeled trees, the increasing, unary-binary, and full binary predicates, and the level-order reading word, and proves the equivalence between this tree-level formulation and the level-matching formulation of Proposition 2.1. The final statements are therefore about genuine trees, with no encoding left to trust:

$$\texttt{theoremA\_trees} : (\exists\, t,\ \mathrm{Inc}\ t \wedge \mathrm{UB}\ t \wedge \mathrm{levelOrder}\,[t] = w) \leftrightarrow \mathrm{HeapCond}\ w$$

$$\texttt{theoremB\_trees} : (\exists\, t,\ \mathrm{Inc}\ t \wedge \mathrm{UB}\ t \wedge \mathrm{levelOrder}\,[t] = w) \to (\exists\, t,\ \mathrm{Inc}\ t \wedge \mathrm{FB}\ t \wedge \mathrm{levelOrder}\,[t] = w)$$

under the hypotheses of the corresponding theorems. The development contains no `sorry`, and the kernel's axiom audit reports only propositional extensionality, choice, and quotient soundness, the standard axiomatic base of Lean's mathematical library. Formalization also repaid the effort mathematically: the original proof of Theorem 3.2 required a delicate case analysis forcing dyadic level sizes, which the formal proof replaced by the telescoping invariant $s_j \leq 2s_{j-1}$ appearing in Section 3, and the explicit chain-shifting arguments first used for Lemmas 4.3 and 4.4 dissolved into the structural inductions printed above.

Data. The twenty-eight new terms, the raw outputs of both programs, and the 47 certificates of Proposition 5.1 are archived in the repository, and the entire verification, computational and formal, reproduces from two commands.

7 Conclusion

The breadth-first reading words of increasing trees looked, from the 2014 data, like a family of six parallel enumeration problems distinguished only by the choice of pattern. The results of this paper show that the landscape is sharply differentiated. For 312 the problem is rigid: realizability collapses onto the single complete shape, the word class is characterized by the elementary heap inequality of Theorem 3.2, and the sequence `A245899` is governed by a Catalan convolution. For 231 the shapes remain genuinely diverse, but Theorem 4.5 shows that at odd length the unary-binary and full binary classes coincide, so the six sequences of [7] contain only four independent enumeration problems. For 321 neither statement survives, and the divergence begins at exactly the first term beyond the published data, a reminder that agreement of small initial segments is weak evidence in this territory.

Beyond the specific identities, two products of this work seem likely to be useful elsewhere. The promotion cascade of Lemma 4.1 is a pattern-free exchange principle for breadth-first realizability, and it did all of the structural work in Section 4; we expect it to transfer to other reading-word problems essentially unchanged. And the three-layer methodology, exhaustive falsification of candidate lemmas before proof, followed by formalization that fed simplifications back into the informal arguments, materially shaped the final mathematics: the printed proofs of Theorem 3.2 and Lemmas 4.3 and 4.4 are shorter and cleaner than the versions first found, and three plausible intermediate statements were eliminated by machine before they could mislead the development. The complete resolution of the conjectures, from open OEIS annotations to kernel-checked theorems about inductive trees, closes every gap between claim and certainty that we know how to close.

8 Open problems

Question 8.1. Find a closed formula or a proven recurrence for `A245898`, the 231-avoiding unary-binary reading words. The sequence matches nothing previously in the OEIS, and its growth ratio appears to approach a limit near 3.5; even the existence of the limit is open.

Question 8.2. Theorem 3.2 collapses 312-avoiding realizability onto the complete shape. Is there a structural classification, for each pattern τ, of the minimal set of shapes needed to realize all words of $\mathrm{UB}_n(\tau)$? Proposition 5.1 shows that for 321 this set is strictly larger than the full binary shapes; how much larger?

Question 8.3. The promotion cascade holds for arbitrary words and seems to be a basic exchange property of breadth-first realizability. Does it extend to increasing k-ary trees for $k > 2$, and do the corresponding parity collapses hold for the patterns whose forced triples survive, in the sense of Section 4, in the k-ary setting?

Acknowledgments

The research, the proofs, the code, and the Lean formalization in this work were developed in an AI-assisted workflow directed by the author, using Claude (Anthropic). Every result is independently checkable through the verification layers of Section 6, none of which requires trust in that workflow. The author thanks the maintainers of the OEIS, whose careful recording of the 2014 data and conjectures made the questions visible.

References

[1] J. Anderson, E. Egge, M. Riehl, L. Ryan, R. Steinke, and Y. Vaughan, Pattern avoiding linear extensions of rectangular posets, arXiv:1605.06825, 2016.

[2] J. Bettinelli, É. Fusy, C. Mailler, and L. Randazzo, A bijective study of basketball walks, *Séminaire Lotharingien de Combinatoire* **77** (2017), Article B77a.

[3] M. Bóna, *Combinatorics of Permutations*, second edition, Chapman and Hall/CRC, 2012.

[4] C. Defant, Proofs of conjectures about pattern-avoiding linear extensions, *Discrete Mathematics and Theoretical Computer Science* **21**(4) (2019), #16.

[5] S. Kitaev, *Patterns in Permutations and Words*, Springer, 2011.

[6] D. E. Knuth, *The Art of Computer Programming, Volume 1: Fundamental Algorithms*, third edition, Addison-Wesley, 1997.

[7] D. Levin, L. K. Pudwell, M. Riehl, and A. Sandberg, Pattern avoidance in k-ary heaps, *Australasian Journal of Combinatorics* **64**(1) (2016), 120–139.

[8] L. de Moura and S. Ullrich, The Lean 4 theorem prover and programming language, in *Automated Deduction, CADE 28*, Lecture Notes in Computer Science **12699**, Springer, 2021, 625–635.

[9] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, published electronically at `https://oeis.org`, 2026. Sequences `A245898`, `A245899`, `A245900`, `A245901`, `A245902`, `A245903`, `A246747`.